Reflecting Nature

Emilie M. Ø. Strandenæs

Thesis presented for the degree of

MASTER OF PHILOSOPHY

Supervised by Professor Carsten M. Hansen
IFIKK, Faculty of Humanities
University of Oslo
December 2015

Your head will collapse
But there's nothing in it
And you'll ask yourself

Where is my mind?
Where is my mind?
Where is my mind?

The Pixies, *Where is My Mind* (1988)

Emilie M. Ø. Strandenæs

Reflecting Nature

2015

Reflecting Nature

Emilie M. Ø. Strandenæs

http://www.duo.uio.no

Print Reprosentralen, Universitetet i Oslo

Cover illustration: Jaques Linard, *Les Cinq Sens (1638)*

Abstract

By what kind of mechanisms do we perceive the physical world beyond ourselves? What is the relation between perceptual processes and natural evolution? What is the nature of a perceptual state? In this thesis I discuss traditional inference theories of perception. These views have suggested that perceptual processes resemble rational processes of inference or computation. I compare traditional inferentialism with Bayesian perceptual science. In *The Origins of Objectivity* (2010) Tyler Burge argues that perceptual science makes non-trivial use of representational notions. Perceptual states constitutively represent a physical world beyond the individual. I present this view in light of the previous discussion. Finally, I explore whether the perspectival feature of perceptual representation suggests that they are subjective experiences.

Acknowledgements

I am grateful to my supervisor Carsten Hansen for introducing me to many of the topics in this thesis, and for continuous encouragement and inspiration throughout the writing process. Sara K. Vikesdal, Jørgen Dyrstad and Sebastian Watzel have provided elucidating comments to the final drafts of this thesis. I thank Sindre Fjeldstad and Nadia Noorman for valuable discussion. I am grateful to the CSMN for their support. Finally, I thank my family for your patience, kindness and insightful feedback. I dedicate this thesis to Ronja.

Contents

.

Introduction

In what sense does perception reflect a specific environment? By what mechanism do our perceptual systems manage to *account for* the environments that surround us?

In the first chapter of this thesis I present *representationalism* about perception — The view that perceptual states, like beliefs, desires and thoughts, have conditions for accuracy.

In chapter two expand on *inferentialism* in perceptual psychology. Traditionally, this paradigm has been associated with the idea that perceptual systems perform sub-personal operations resembling rational processes. I present Bayesian approaches to perception, which model perception processes as probabilistic computation. I argue that the Bayesian approaches are neutral on the nature of actual perceptual processes.

In *The Origins of Objectivity* (2010) Tyler Burge argues that sub-personal perceptual processes *mirror* physical principles because interaction between individual and environment has shaped the processes. Perceptual systems have developed through a process of natural evolution. The mechanisms that govern perceptual formation reflect the nature they have evolved within, but they do not represent it in computational or inferential processes. In chapter three I present and discuss this view.

While perceptual processes themselves do not represent environmental information, perceptual states are objective representations of the physical world according to Burge (Burge 2010:1). He argues that

perceptual science presupposes that perceptual states can be *accurate* about physical entities. Physical entities figure constitutively in the type individuation of perceptual states in perceptual psychology. The success of perceptual science suggests that representational perceptual states are *real entities.* Perception marks the beginning of representational mind (Burge 2010:xi).

In chapter four I discuss the connection between mental representation and conscious experience. I suggest that the notion of an unconscious mental representation is problematic. How can a perceptual state constitute the perspective of a subject, if not as an experience? Mental representations constitute a *mode of representation*, or *way of referring*. I present a view of such modes of presentation as constitutively subjective in a way that only experiences are.

Ch. 1 Perception

A central paradigm within contemporary philosophy and psychology of perception is that perceptual states can be, and generally are, accurate about entities and subject matter within the physical world. Most notably, *representationalism* is contrasted to direct realist views. Broadly construed, the latter views suggest that entities and subject matter are not *represented* in a perceptual state. Rather, the entities and subject matter are *parts* of the state in question. I make use of central notions from Tyler Burge's (2010) account of representationalism in this section. Representationalism about perception is the view that perceptual states are *intentional* states. Intentionality is the minds directedness upon subject matter. Certain mental states seem to be directed upon subject matter, in the sense that they are *about* that subject matter. Beliefs, desires, thoughts — these are about certain subject matters. There is a significant contemporary tradition of explaining intentionality in terms of *representation.* According to this line of thought, a mental state is about certain subject matter, if it can be *true* about that subject matter. Representationalism about perception is the view that perceptual states represent in this manner. 'Truth' is a predicate generally applied to *propositional contents.* Perceptual states are generally not taken to have propositional contents. *Veridical*[1] perceptual states are *accurate* (Burge 2010:39). Perceptual states have perceptual contents with accuracy conditions, according to representationalism. Veridical perceptual states accurately indicate the environment to individuals (Burge 2010:39). *Mental representation,* intentionality and representation are used interchangeably in this thesis, unless something else is indicated.

[1] 'All kinds of being 'true', 'correct', 'accurate' etc. are sub-cases of veridicality.' (Burge 2010:39)

1.1 Objectivity

In his major work *The Origins of Objectivity* (2010) Tyler Burge argues that representationalism is assumed by perceptual psychology. The science as it is today relies on the assumption that perceptual states are representational. Representational vocabulary not only figures in perceptual science — it plays a non-trivial role in psychological explanation (Burge 2010:292). Perceptual science provides extensive reason to think that perceptual states are *mental representations.* In fact, it suggests that mental representation *begins* in perception — phylogenetically and for the individual (Burge 2010:xi). Intentionality *marks* the mental. *Perception* is the beginning of the mind (Burge 2010:xii).

Objectivity is a notion that figures centrally in Burge's account of representationalism. As mentioned, intentionality is the minds *directedness upon subject matter.* Intentionality *begins* in perception in the sense that it is the 'most elementary type of representation' (Burge 2010:xi) namely, 'accurate — objective — representation of the physical world' (Burge 2010:xi). Perceptual states are directed upon subject matter in a special sense: They are directed upon the physical world and they are (generally) accurate about that world. This way of being directed upon the world is the sense in which perceptual states *objectively represent.* Two main features are associated with objectivity: Representation of physical reality, and *accuracy* (Burge 2010:46). Perception just *is* the individual's capacity to '(…) represent the world objectively' (Burge 2010:1).

'The objectivity of such representation lies (…) in it's accuracy and it's specifying attributes relevant attributes in a way that entail their physicality. It also lies in the physical subject matter's being (…) constitutively non-perspectival'. (Burge 2010:59)

Objectivity is intended to capture the phenomenon that individuals have mental states about physical reality beyond themselves. These states are generally accurate. Perception generally informs perceives about how things *really are* with that physical reality. Perceptual states manage to reveal something *beyond* that perceivers' idiosyncratic 'point of view.' As perceptual states are *representational,* they will always be *perspectival.* Representational states have *modes of representing* — they have *representational content* (Burge 2010:37). Representational contents '(…) constitute, or help constitute, modes in which an individual thinks about, intends or perceives a subject matter.' (Burge 2010:38) Hence, perceptual states have *ways* or *modes* in which subject matter is represented. The individual perceives reality beyond herself from a *perspective,* or point of view. However, perceptual states are about subject matter that is '(…) constitutively non-perspectival.' (Burge 2010:59) Hence, in perception, an individual has awareness of an *objective, non-perspectival* subject matter, from her *subjective, perspectival* point of view.

Understanding how individuals can have representational states about an external physical reality is a major philosophical problem.[2] In what sense are our minds about physical reality? How can our subjective perspectival mental states come to be accurate about the world beyond ourselves? Burge argues that veridical representation of physical reality begins in perception (Burge 2010:23). Hence our sense perception underlies other kinds of mental representation of an external world. A central problem associated with perception as objective representation is the underdetermination of sensation. The problem expresses an idea that has been prevalent within philosophy. How does the information we register in our sensory organs result in accurate physical

[2] '(…) how to combine the perspective of a particular person inside the world with an objective view of that same worlds, the person and his viewpoint included?' (Nagel 1986:3)

representation? How is perception objective representation of something beyond our sensory organs?

1.2 The Problem of Underdetermination

Understanding the relation between sensation and *perceptual constancy* is one of the key themes in modern perceptual psychology. I expand on perceptual constancy in chapters two and three. For now, I think it is sufficient to say that *objectivity* is what perceptual constancy *provides* to perceivers. Burge argues that '*The primary problem for the psychology of visual perception'* (Burge 2010:89) is in fact to explain how visual perceptual states that objectively represent the physical environment 'are formed from the immediate effects of proximal stimulation (…)' (Burge 2010:89). The proximal stimulation that our sensory organs register underdetermines their environmental causes. Hence, sensation does not in it self reveal how perception can be objective representation of a *non-perspectival* physical reality. Proximal sensory registration does not in it self appear to determinately reveal a *non-perspectival* physical reality.

1.3 Underdetermination in Vision

To explain the problem of underdetermination in modern perceptual psychology, I will look to the explanation of the problem in Wade & Swanston (1991).

All sensory systems function by transuding some type of environmental energy into a form that can be analysed by the cells in the central nervous system (Wade and Swanston 1991:59).

When we sense, nerve cells in sensory organs are influenced by environmental energy[3]. These cells stimulate other cells until signals finally reach the CNS (Central Nervous System). Nerve cells, or neurons, influence one another by transmitting electro-chemical signals. Cells transmit signals that influence other cells across synapses. In vision, the retina is where environmental energy is registered by neurons. On the retina, there are receptors that contain light-sensitive pigments. The retina is an outgrowth of the CNS. When light bounces off environmental entities and hits the eye, the light rays are 'concentrated' by the lens and the cornea. Light is transmitted through the pupil, and reaches the retina. It casts an inverted image of the environmental scene that it was reflected from. The image cast on the retina is a 2D representation of a 3D environmental scene.

The environmental cause of visual stimulus is underdetermined. A 2D image of a 3D scene does not carry information about how it should be interpreted as representing a 3D scene. Many different environmental scenes could have caused the 2D image that is cast on the retina. The nature of retinal sensory registration entails that the distal environmental cause of proximal visual stimulus is underdetermined. Hence, the 2D visual stimulation that underlies visual perception is not alone sufficient for visual perception.

In our external environment, the strength and location of light rays shift. Shades move over surfaces that we perceive, sunrays become weaker as clouds pass over the sky. Variances in lighting conditions and motion influence retinal images. When I view a bicycle in motion, the size of the retinal image cast by the bicycle will change. When I walk toward my friend, the image he casts will be larger as I approach. Visual stimulus

[3] The following account is, as mentioned, based on the explanation of *visual underdetermination* in Wade & Swanston (1991:55-65).

drastically varies. Yet, constant physical properties are revealed in vision. I perceive my friend as remaining constant in size even though there is relative motion between us. It seems to me that a field of oats remains constant in colour although a windy day makes cloud-shaped shades move rapidly over the field.

Some psychologists and philosophers have maintained the view that perceivers must somehow internally *interpret sensory data* in order to have determinate perception. Perceivers must entertain some capacity to *take* sensory information to be about a specific environmental cause. A dominant tradition within this paradigm is *inferentialism*. Views of this kind suggest that perceivers, or perceptual systems, must somehow make inference, from sensory information and some additional information, to perceptual representation. In the following chapter I discuss two central examples of inferentialist theories.

Irvin Rock is a key figure in the development of modern inference theory, and I present and discuss his notion of perceptual inference. I go on to present Bayesian perceptual science. This approach models perceptual processes as probabilistic computations. While Bayesian perceptual science resembles inference theory, it does not suggest that actual perceptual systems perform these computations. Hence, the Bayesian approach is as *inferentialist approach* only in so far as it models perception as computation. It differs from traditional inferentialism because it is silent on the nature of actual perceptual processes.

The science I present reveals how perceptual processes must in some sense reflect principles that govern the environments of actual perceivers. Principles of physical reality are somehow *evident* in perceptual processes. While Rock argues that these principles are

represented within perceptual systems, Bayesian models are silent on just how actual perceptual processes *mirror* natural environments.

In 2010, Burge argues that, while environmental principles are *reflected* in perceptual processes:

> (...) there is no sense in which the principles are "accessible" to the perceiver or the perceiver's perceptual system. (Burge 2010:96)

Hence, Burge's view rejects the idea that actual perceptual systems make inference or perform computation. I contrast Burge's view with inferentialism in chapter three. Assuming that perceptual systems take sensory information to *mean* something, is '(...) almost as bad as thinking of the planetary system as applying principles governing its motion.' (Burge 2010:96)

Ch. 2 Perceptual Psychology

I order to introduce inferentialist theories I will expand on perceptual constancy. I will rely on Rock's own description of size constancy, as I think his account presents a good vantage point for the general discussion of inferentialism.

2.1 Size Constancy

In the chapter on size constancy from his book *Introduction to Perception* Irvin Rock addresses two problems regarding size perception (Rock 1975:27). One problem concerns how the size of objects appears constant to us, even as entities move closer and further away from us. Why do objects not appear to decrease and increase in size as they move? Images cast on the retina through the ocular lens do. This question addresses just *how* things can appear[4] to have constant size. The other question Rock addresses concerns the relation between visual stimuli and the perceived size of objects: Why do things appear to have the specific sizes that they do? The two questions relate to the same phenomenon: Thing appear to have stable and specific size. This phenomenon is size constancy.

Both of the questions Rock addresses make methodological use of the problem of underdetermination. Rock presents a version of the problem by giving a description of vision in terms of optics and a description of the eye. The size of a retinal image is inversely proportional to the distance of an object: Not only is the size of the retinal image

[4] Rock makes no terminological distinction between 'appearance' and 'perception'. Hence, he seems to assume that perception *is* perceptual experience.

proportional to the size of the object that reflects light onto the ocular lens: But as the ocular lens projects an inverted image of the physical scene onto the retina, the size of the inverted image will be proportional to the angles with which light hits it.

The following illustrations are taken from (1975:28-30) Rock explains how size constancy is affected by distance. The distance between the eye of the viewer and the object viewed, influences the size of the image on the viewer's retina. (Rock 1975:28) The eye contains a lens. A lens brings rays of light to focus. It bends incoming rays of light, in such a way that the light that is transmitted from — or reflected from — a point in space, if allowed through the pupillary opening of the eye, will be focused on a point in the eye. Hence, points in space have corresponding points on the retina. Light from point A in space will create retinal image-point a.

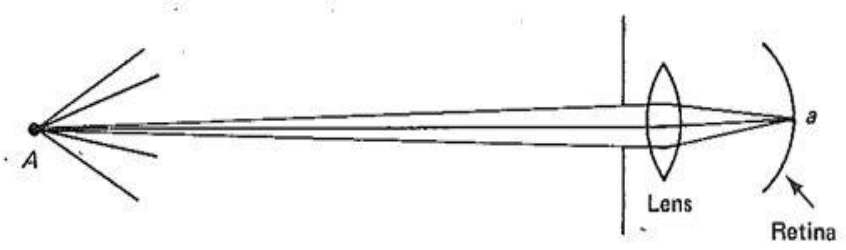

Figure 2-1

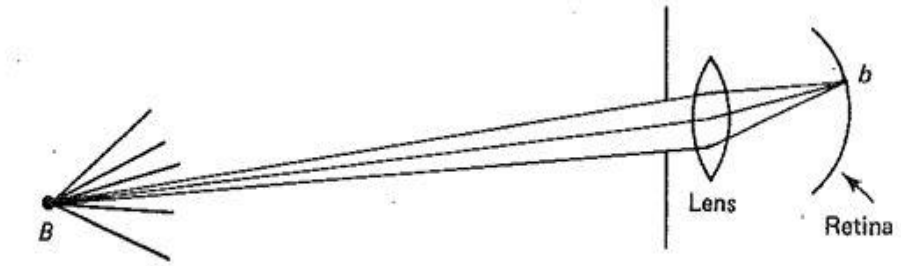

Figure 2-2

Imagine that points A and B, as illustrated in the figures are the top and bottom of a physical object. The further apart these two points are, i.e., the larger the object is, the larger the retinal image the object casts will be. This point is simply an environmental fact. A larger distance between A and B causes a larger distance between a and b. This is because the direction of the light-rays determines the direction in which the ocular lens will deflect them, as seen in the illustrations below (Rock 1975:29).

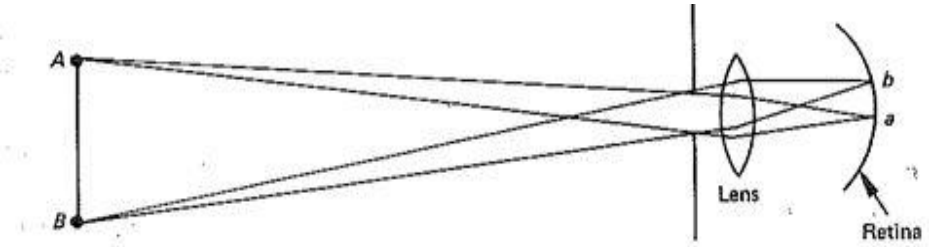

Figure 2-3

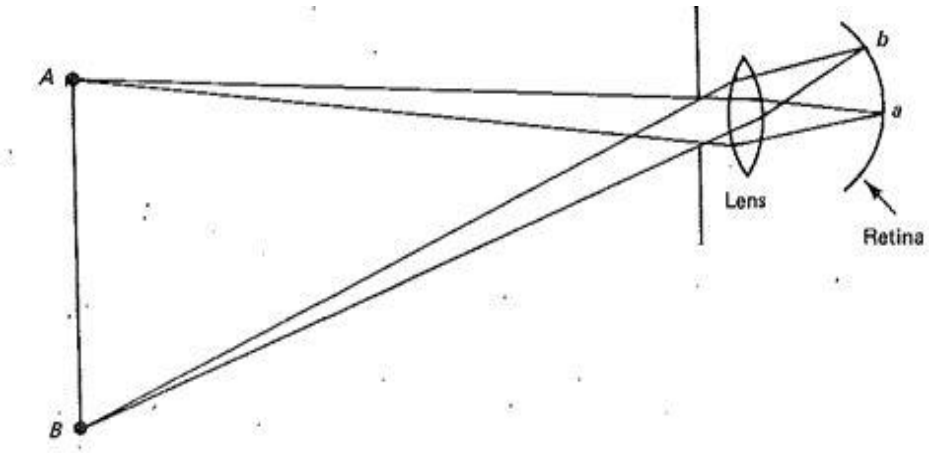

Figure 2-4

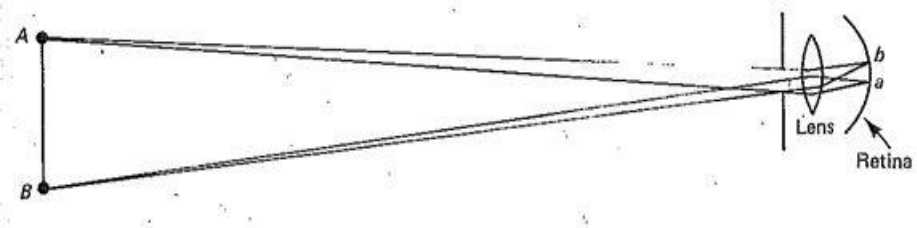

Figure 2–5

Should the object A-B be moved further away from the ocular lens, the difference in the direction of light-rays from point A and B would be smaller. Hence, light rays from these points would deflect closer together — points a and b would be closer (Rock 1975:30). Hence, distance affects the size of retinal image. As Rock's puts it:

(…) the size of the retinal image (or visual angle) is inversely proportional to the distance of the object. The term *visual angle* is used synonymously with size of retinal image (…) Fig. 2-6. (Rock 1975:30)

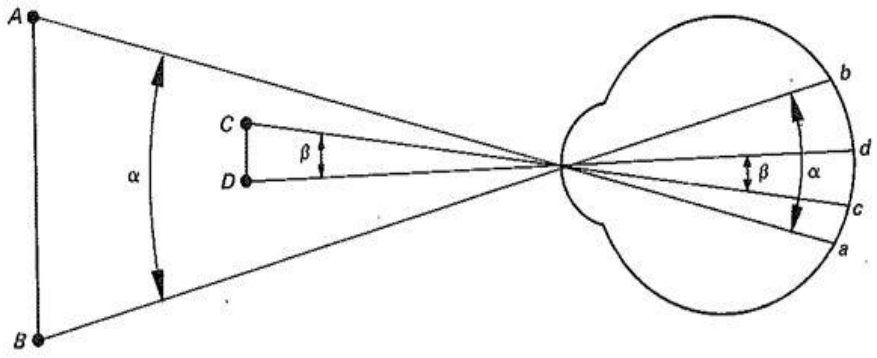

Figure 2–6

These illustrations reveal how the angle between points of registration on the retina increases proportionally to the angle between the points in

space that emit the light registered. The angle between points that emanate light that is registered on the retina decreases relatively to the viewer as the points move further from the ocular lens.

Rock's examples show that the size of a retinal image is simply a function of the size and distance of an object looked upon. However, the perceived size of objects appears not to be a result of the same function as the size of retinal images (Rock 1975:31).

Rock points out how the appearance of size in perception *can be deceitful.* In a sense, objects do appear smaller at a distance (Rock 1975:32). However they do not appear *proportionally* smaller to the increased distance between viewer and entity viewed. Things perceived at a distance don't look as much smaller as the retinal image decreases with distance. The size of retinal images varies more radically than the environment appears to us to do (Rock 1975:31).

The illustrations above reveal that visual angle is ambiguous regarding the size of objects. Hence, in order to have determinate perception of the size of objects, the distance between perceiver and object must somehow influence the relevant perceptual process.

Rock's account of size constancy shows that certain physical facts are somehow reflected in perceptual processes: Size constancy somehow reflects how distance affects the angles of light-rays that enter the pupillary opening.

So far, the account I have given of Rock's explanation of size constancy does not presuppose that perceptual processes are inferences. However, Rock explicitly claims that distance is '*taken into account*' in size constancy (Rock 1975:33). Size is evaluated by the perceptual system

(Rock 1975:75). Hence, the relevant physical facts (physical facts that explain how distance affects retinal image size), are reflected in the perceptual system in the sense that they are represented in it.

2.2 Reflecting as Inferring

In the article 'In Defence of Unconscious Inference' (1975b), Rock makes explicitly what he means by unconscious inference.

> By unconscious inference I mean that the process of arriving at the percept is one much like reasoning in which conclusions are drawn from premises, except that in perception the outcome is a percept rather than a conclusion. (Rock 1975b: 258)

He suggests that facts from optics and physics (facts regarding the way distance affects retinal image size) are stored in the perceptual system as premises. These premises figure in perceptual processes that resemble the process of reasoning. Principles that determine how sensory stimulus and physical facts will yield objective representations are not merely describable for the perceptual system on this view. Rather, the perceptual system makes use of principles of reasoning. Size constancy is (something like) a *judgment,* starting from environmental stimulus (retinal image) and environmental principles (facts of optics and physics) and ending at stable perceptual representation. Rock commits to the view that actual perceptual systems make inferences.

The problems involved with the notion of unconscious inference have been subject to extensive philosophical discussion[5]. In 'Perception as Unconscious Inference' Gary Hatfield discusses three central and intuitive problems associated with such a notion: *The Cognitive Machinery problem* expresses the concern that theories of unconscious

[5] See Ludwig (1996) and Burge (2010:92-93)

inference must 'account for the cognitive resources needed to carry them out'. Hatfield questions whether unconscious inferences are supposed to be inferences made by the cognitive faculty *of an individual,* or whether the perceptual system itself is supposed to entertain the capacity to perform inference (Hatfield 2002:120). *The Sophisticated Content Problem* questions how the perceptual system or perceivers in general are supposed to entertain premises regarding their environments. For example, if size constancy requires that perceptual systems or perceivers somehow entertain *premises* regarding how distance affects retinal image size (note that this information plays an essential role in Rock's explanation of size constancy), how can such premises be entertained by visual systems or perceivers? It is not the case that individuals need to understand the principles that govern size constancy in order to perceive entities as having specific, determinate sizes. If size constancy is a matter of inference, must the visual system understand these principles? (Hatfield 2002:120)[6] *The Phenomenal Experience Problem* questions how the conclusion of an inference can be anything like a perceptual experience. How can premises containing environmental information about distance and size result in a *conclusion* that is *visual experience* of objects? (Hatfield 2002:120) These problems are just examples of the kind of issues facing theories of unconscious inference.

Rock's theory has explanatory value that is independent of the notion of unconscious inference. His account of size perception reveals that distance will somehow be reflected in size constancy. Even if one denies that the perceptual system evaluates distance, it is explanatorily interesting that size constancy somehow relates to distance.

[6] Both of these problems essentially illustrate Burge's arguments against Compensatory Individual Representationalism: They illustrate how certain theories require inappropriate capacities of perceiving individuals (Burge 2010: 13-22, ch. 4-7).

I turn to Bayesian perceptual science. As mentioned, this approach does model perception as an unconscious probabilistic computation. However, the Bayesian approach does not commit to the view that actual perceptual systems perform such computations.

2.3 Bayesian Approaches

Bayesian perceptual science describes how perceptual systems reflect environmental principles, by creating probabilistic, mathematical models. Bayesian approaches model perceptual processes in the following way: The perceptual system entertains a hypothetical space of possible environmental scenes. It selects the most *likely* scene based on input in order to overcome the problem of underdetermination. Hence Bayesian perceptual psychology rests on the idea that certain hypotheses about the relation between a proximal stimulus and an environmental cause have a higher likelihood than others. Bayesian statistics are based on Bayes theorem. The following example can illustrate the theorem, applied to a case where posterior statistical data informs a prior likelihood.

A team of scientists are mapping the occurrence of a rare, dangerous disease in a population. They want to establish the frequency of affected individuals within the population. Estimates suggest that 2‰ of the population are afflicted. (a) represents the frequency of cases within the general population. In order to map the disease, the scientists conduct genetic screenings of all known patients. They discover that 60% of patients have a specific genome. The relevant genome has been given much previous attention and is well mapped: Approximately 1% of individuals within the general population have the genome. (b) represents the value of the distribution of the genome in the general population. With the present data the scientists can use Bayes' theorem

in order to calculate the likelihood of individual being afflicted (Pa) if they have the genome (Pb). They can express how the prior likelihood of any individual being afflicted (a) is informed by additional relevant statistical data, namely the data they have regarding the relation between the disease and the genome (Pb|a) and the data they have regarding the general distribution of the genome (b).

Bayes' theorem is as follows:

$$(Pa|b) = \frac{(Pb|a)(Pa)}{Pb}$$

We know the value of (a), (b) and (Pb|a).

$$(Pa|b) = \frac{(0.6)(0.002)}{0.01} = 0.12$$

The happy team of scientists can publish a break through article establishing that people with the relevant genome have a 12 % chance of being afflicted by the disease, and receive funding for further research.

Bayesian *decision theory* is based on Bayesian statistics. This is a useful as a tool for modelling cases where several hypotheses are underdetermined by data. As mentioned, Bayes theorem provides a tool for calculating how a prior likelihood will be informed by additional relevant data. Hence, Bayesian decision theory can be used to model cases where the likelihood of a given hypothesis is informed by additional testing. The example above illustrates this. Bayesian decision theory effectively models how a given hypothesis has a higher likelihood of being the solution to an *underdetermination* problem considered in light of some additional relevant data.

Bayesian decision theory is apt for describing mathematical models of perceptual processes. The models can be used to describe how perceptual systems overcome the problem of underdetermination by reflecting the statistical properties of the environments they have developed in. I will present a simplified and generalized example to illustrate this point. It is a statistical fact about our natural environments that sources of light are generally placed above us. Hence, one can describe a model of a perceptual system where different possible locations of light sources are represented as hypotheses about the environment in the perceptual system. Each hypothesis will have an equal prior probability. Interaction between perceivers and the environment will influence the probability each of the hypotheses. As light generally *does* come from above, this hypothesis will be increasingly probable as perceivers and environments interact. Hence, Bayesian models can incorporate how interaction with an environment will make specific hypotheses about that environment more probable. They model the perceptual system as probabilistic inference. Specific entities and subject matter will be represented in perception if they are the content of the 'hypothesis' that is the most likely. On this simplified picture, the environmental hypothesis with the highest likelihood will be the one that is recurrently confirmed by the actual environment. The models can incorporate how interaction with an environment makes specific hypotheses about that environment more plausible.

Bayesian approaches to perception do not presuppose that the interaction between a specific individual and her environment informs the likelihood of environmental hypotheses. Many of these models suggest that interaction between species and environments *over time* is reflected in perceptual processes. Such models presuppose that actual perceptual processes reflect that perceptual systems have been selected for (Geisler & Diehl 2003:379).

Actual perceptual processes will mirror what has been beneficial to perceive under which environmental circumstances for individuals' ancestors. This is evident considering the use of *ideal observers.* An ideal observer is a Bayesian probabilistic computation that represents an *ideal solution* to a perceptual task (Geisler & Diehl 2003:385-387; Geisler 2011:771-772). An ideal solution to a perceptual task is modelled in light of meta-data regarding the *utility* of predicting a given environmental cause under certain environmental conditions. Ideal observers incorporate *utility:* They recognize which perceptual outcome is useful under what environmental conditions.

As mentioned ideal observers illustrate a central feature of Bayesian approaches, namely that they attempt to capture how perceptual systems developed by a process of natural evolution. Ideal observers rely on the idea that actual perceptual systems will reflect *utility:* Actual perceptual processes will be determined by evolution. Hence individuals with perceptual systems that indicate the environment in *such a way* that the interests of those individuals are well preserved will be selected for. Ideal observers not only specify the likelihood of an environmental cause given a stimulus, they predict the likelihood of a given perceptual prediction, given the probability of an environmental cause, and a meta-consideration of the utility of predicting in such a way under given environmental conditions (in the presence of given physical entities/states of affairs). An ideal observer can be a good model for describing perceptual phenomena where properties or entities that do not exist in the environment are repeatedly indicated to individuals. An example, of such misperception might be how rabbits frequently misperceive predators in cases where no actual predator prevails. Rabbits frequently thump their feet in order to communicate to other rabbits that danger is approaching, when they hear loud noise or observe

rapid, unexpected movements. Perhaps this illustrates that the gain of being alert is larger than the cost of misperceiving in certain cases. Ideal observers can account for such phenomena.

Modelling the process from sensory input to output as a probabilistic inference allows statistical analyses of the utility of specific outcomes under specific environmental conditions. This is illustrated by the values that must be specified for a description of an ideal observer.

There are four values that are specified in an ideal observer: (1) A perceptual task; (2) the prior probability of a category (an environmental state); (3) the likelihood of any possible stimulus given each of the possible categories (each possible environmental state); (4) a utility function (the costs/benefits associated with predicting a given state) (Geisler & Diehl 2003:380).

Initially, ideal observers were specified for simple models of perceptual tasks: They were typically set to solve tasks of accurately predicting between two alternative candidate environmental states. For such a task, specifying the utility function is simple: There is an equal cost involved with any erroneous prediction. An example of such a task might be to accurately detect a spot of light against a white noise background. This task involves selecting between two possible alternatives at a given time (spot or no spot) (Geisler & Diehl 2003:380). Describing an ideal observer for this task involves describing the prior likelihood of every possible state of the environment (the likelihood of there being a light spot at time t1; at time t2; etc.); The likelihood of every possible stimulus being caused by every possible state of the environment (light intensity l1 caused by light spot at time t1; etc) and a value for the utility function (the cost/benefit of reaction r (the formation of a given state) if there is a light spot at t1, etc). Perceptual systems have

developed through natural selection, so descriptions of the process of perceptual prediction must reflect what behaviour is beneficial for an organism under which environmental conditions: It must reflect that individuals who make beneficial predictions are selected for (Geisler & Diehl 2003:379).

Natural tasks are vastly more complicated than the simple tasks that were initially described in ideal observers. Natural tasks are not simply about maximizing the accuracy of prediction: In describing ideal observers for complex natural tasks, this has to be incorporated (Geisler & Diehl 2003:381). Understanding how actual populations of species have developed informs how the scientist think about the value of the utility function for ideal observers: Investigating what actual conditions correlate with the development of which perceptual capacities provide insight into when it has been useful for individuals to react in which way to a stimulation (Geisler & Diehl 2003:381).

For many advanced computations of solutions to advanced natural tasks, Bayesian inferences that specify the utility function as *maximizing fitness* have provided results that largely correspond to actual organism's predictions. Maximal fitness observers are ideal observers where the utility function represents *statistical data* regarding what reactions *have actually proven beneficial* for organisms under what environmental conditions (Geisler & Diehl 2003:381). This suggests that actual perceptual systems do reflect the utility of certain reactions, given certain environmental causes. Indeed, a maximal fitness observer that is set to solve a task of detecting contours based on information equivalent to sufficient visual stimulatio, yields predictions that are nearly equivalent to what humans do on the basis of visual stimulation (Geisler & Diehl 2003:396). The upshot of this is that human vision appears to illustrate that there are cases in which accurate perception is beneficial.

Some human visual perception is accurate because accurate vision has been beneficial for survival. I address this idea in relation to Burge's picture of teleology and accuracy in section 3.6.

2.4 Probabilistic Perception

Bayesian approaches assume that the perceptual system must somehow be 'tuned in' on an environment. Just as distance must be *reflected,* or *play some role* in size constancy, the likelihood of environmental scenes given stimuli, and the utility of indicating certain entities in specific situations, must be reflected in perceptual processes, according to Bayesian models.

How does the system reflect likelihood and utility? There are two ways to interpret Bayesian perceptual science: Perceptual systems can operate on non-computational mechanisms that can be modelled as probabilistic inferences, or they can represent the probabilities and compute on them. As mentioned, there is no reason to think that Bayesian approaches are committed to anything stronger than the former view.

The first interpretation involves understanding perceptual process as some kind of transitional process from stimulus to state that can be described as values in a Bayesian function. Actual perceptual systems do not compute, on this view: However, perceptual processes are computable. Both the probability of environmental cause given state and utility would be reflected in perceptual processes, in the sense that one cold observe that these factors when studying perceptual processes. One can describe computational models for the weather or social economic structures: These phenomenon do not perform computations. In the same way, perceptual systems do not perform computations even though they can be computed on, on this view. The fact that Bayesian decision

theory is used to describe perceptual processes does not suggest that perceptual systems make decisions: Only that they must overcome an underdetermination problem.

One can interpret perceptual systems as representing probabilities and performing computations. The system would normatively follow laws that can be described as values in a Bayesian function. It would literally perform computations on the basis of these functions. On this interpretation, the system would represent statistical facts about the environment and the relation between the environment and stimulation, in order to calculate the probability of an environmental cause, based on stimulation. Independently of the exact process of formation, on this understanding of the perceptual system, it would not only be mathematically useful to describe the formation of perceptual states as an inferential operation: Rather, this would describe the actual process of formation.

Bayesian decision theory is used in any number of sciences. Social medicine, social studies and natural science make use of Bayesian statistics and Bayesian decision theory. The fact processes can be modelled using probabilistic decision theory does not in it self suggest that the process involves any decision-making.

Explicit statements from psychologists working with Bayesian models vindicate the first interpretation of the relation between Bayesian computations and actual perceptual systems:

'(…) the terms in the formulas of the Bayesian framework can represent any psychological/biological system that can be characterized by an input and an output.' (Geisler & Diehl 2003:399)

This quote essentially reveals that psychologists working within the Bayesian framework do not commit to the view that perceptual processes are computations. Only in so far as there are reasons for taking any 'biological system that can be characterized by an input and output' to perform computations, are there reasons for assuming that actual perceptual systems do. Bayesian approaches to perception are in them selves neutral on the *inferential nature* of actual perceptual systems. The models themselves do not suggest that perceptual systems are computational.

2.5 Inference and Representationalism

Rock's notion of perception as unconscious inference commits him to representationalism about perception. Suggesting that perceptual states are the conclusions of processes of *reasoning over premises,* entails that perceptual states have representational content. They *follow as consequences* from represented premises and rules. Hence the states must be representational states. Any view suggesting that actual perceptual processes are inferences or computations with perceptual states as their conclusions or outcomes simply claims that perceptual processes are representational processes with representational outcomes.

As I have argued, Bayesian models do not address whether actual perceptual processes are computational. Hence, they are not obviously committed to representationalism for *the same reasons* as other inference theories. However, there is reason to think that the theories do assume representationalism.

Bayesian approaches individuate perceptual states by describing them as representations of specific environmental entities. They describe perceptual states as states that can be accurate about particular objects

or properties in the environment:

> However, as a concrete example, consider a task where there are just two categories of object and the observer's (authors remark: the *ideal* 'observer's', i.e. the computational model's) task is to be as accurate as possible in identifying which object was presented.' (Geisler 2011:772, italics mine)

> In general, it is true that much of human perception is *veridical* under natural conditions. (Geisler & Diehl 2003:397, italics mine)

While Bayesian models do not assume that actual perceptual systems perform computations, they do individuate perceptual states *as representational contents with accuracy conditions.* They assume that the outcomes of perceptual processes are perceptual states about physical entities. Perceptual states are assumed to have representational content *in the models.* The solutions to perceptual tasks are contents about specific physical entities[7].

This does not entail that the use of such representational notions reveals that actual perceptual processes have representational outcomes. Modelling perceptual states as probabilistic inference relies on specifying perceptual tasks that have representational contents as their solution. A perceptual state analysed *as the outcome* of a computation will naturally be a representational content *in that model.* In this sense, the theories might be analysed as neutral on whether perceptual states genuinely have representational contents.

A central question is how a non-inferential, non-computational process can result in perceptual representations. In (2010) Burge explicitly

[7] This point is recognized by Michael Rescorla in 'Bayesian Perceptual Psychology' (2013:14): 'Bayesian models individuate both explananda and explanantia in representational terms.'

argues that no pre-perceptual representation is required for objective veridical representation. Hence he argues that representationalism is true, without supposing that perceptual systems represent *conditions for representing*. Perceptual processes involve no representation of conditions for objective representation of the physical world (Burge 2010:19). Representation of conditions for objective representation is precisely what Rock requires from perceptual systems. Perceptual systems, according to Rock, have to represent environmental facts in order to make inference from sensory data to perceptual state. Evidently, Burge's reasons for thinking that perceptual states are representational are not the kind of reason I described in relation to traditional inference theory.

Ch. 3 Origins of Objectivity

Perceptual processes somehow reflect principles that govern the environments around us. They mirror statistical principles that are prevalent in the physical world. This is *required for* perceptual constancy. In (2010) Burge argues that perceptual processes do not *represent* these principles. How do the systems reflect them?

As evident from Bayesian perceptual science, evolution shapes perceptual processes. Somehow, natural evolution plays a significant role in determining the mechanisms of perception. What is the relation between *accuracy* and evolution? Burge argues that perceptual states objectively represent. They are generally *accurate* about the physical world. In what sense are processes shaped by natural evolution aimed at accuracy?

In order to explain Burge's view I present his accounts of *lightness constancy* and *planar slant/planar surface texture*. These examples reveal that perceptual processes mirror statistical regularities in the environment. In light of these I discuss the teleological element of his picture.

3.1 Reflecting Regularities

Lightness Constancy

Lightness constancy is the capacity to visually perceive a surface as the having the same lightness (the same colour) even as differences in the illumination of the surface provide an organism with drastically varying

proximal stimulation (Burge 2010:351). Many organisms are generally capable of perceiving achromatic surfaces as having an even colour under varying conditions of illumination. A surface seems uniformly white, even if parts of the surface are shaded. This means that organisms are capable of keeping track of the *surface reflectance* of an entity.

The intensity of light that a surface reflects is given by a combination of the reflectance properties of the surface material and the intensity of light that is directed at the surface (the illumination of the surface). This is a fact about the physical environment. A surface reflects some intensity of light. The intensity of the light varies as the illumination of the surface varies. The receptors that register light intensity cannot in themselves determine whether lightness intensity variations are due to variations in surface reflectance or in luminance (Burge 2010:352). This provides a clear example of how proximal stimulus alone does not privilege one representational perceptual state over another. It illustrates underdetermination.

There is nothing about the light intensity registered by the receptors in the eye *alone* that should indicate that a white surface with shades appears the same shade of white all over, rather than appearing patterned in different shades of white and grey. Visual psychology provides experiments to the effect that the capacity to make such a distinction is not a higher cognitive capacity. Many primitive animals are able to distinguish surface reflectance from surface illumination. They have the capacity of perceiving lightness constancy. This indicates that in determining lightness constancy the *perceptual system* must draw the distinction between what is a property of a distal object and what is an environmental condition that, in effect, distorts visual stimulation from the object (Burge 2010:352).

Appealing to, amongst other, vision scientist Steven Palmer, Burge claims that the receptors that register the intensity of light reflected from a surface *immediately* form what are called luminance contours, on the basis of discontinuities in the light intensity registered by adjacent receptors in the eye (Burge 2010:352). Luminance contours are functional registrations of discontinuities in light intensity. If two neighbouring receptors register sufficiently different intensities of light, corresponding to there being some sudden difference of light-intensity in the perceived environment (as there is if one observes a white surface with dark patches) the receptors register a luminance contour. The registration of luminance contours indicates that there is some sort of discontinuity of lightness intensity in the distal environment: The luminance contours that the receptors register correspond to some kind of environmental discontinuity.

> If a series of spatially adjacent receptors (…) produce a pattern of registrations of sharply different, adjacent levels of light intensity, (the) receptors produce a luminance contour. (Burge 2010:352)

Lightness contours are direct, non-perceptual encodings of proximal stimulations. The formation of lightness contours alone still underdetermines the distal cause of the contours. There is nothing about the contours themselves that indicate whether they are caused by discontinuities in surface reflectance or discontinuities in luminance. If the visual system is to overcome this problem, something must determine that a visual perceptual state is an appropriate reaction to a given luminance contour.

On Burge's account, law-like regularities between states of the perceptual system (Burge 2010:346) do this work. In lightness constancy the perceptual system operates on semi-automatic transitions from stimulus to luminance contours to states indicating edges. These

three steps reflect three stages in the process of objective perceptual representation. The step from sensory input; to registration; to perceptual indication, or representation.

'Law-like regularities in the perceptual system (…) reflect (…) law-like regularities in the distal environment.' (Burge 2010:346) Perceptual processes have developed so form perceptual constancies in accordance with statistical regularities in the environment. This reveals a deep similarity between Burge's account and Bayesian models:

> Perceptual systems have developed so that their representational states tend to correlate with the likely causal antecedent, in the systems' formative environment, of the given proximal stimulation. There is a many-one mapping from distal environmental cause, to the proximal stimulus, and a one-many mapping from proximal stimulus to the environment. But there is something like a one-one mapping from proximal stimulus to distal environmental cause that is most likely to have generated that proximal stimulus. (…) Nature molds all sensory-systems — perceptual and non-perceptual — to be likely to respond to conditions that are beneficial to animals' function. (Burge 2010:345)

In effect, Burge argues that perceptual systems have developed so as to indicate the most likely environmental cause of a proximal stimulus. However, the system must reflect evolution: It must somehow mirror that natural selections mold systems in a manner such that they are beneficial for the individual's conditions. I address this further in 3.6.

A statistical fact about nature is that sharp discontinuities in the intensity of light reflected by a surface are usually due to discontinuities in surface reflectance, while gradual discontinuities are due to discontinuities in illumination. Registrations of sharp luminance contours will generally result in states indicating a discontinuity in surface reflectance, while gradual luminance contours will generally

result in states indicating a discontinuity in illumination. The registration of a certain luminance contour will result in the formation of a determinate state in a specific case. Hence lightness constancy reflects statistical facts about the natural environment. There is a higher likelihood that a sharp discontinuity in lightness intensity is caused by a surface reflectance discontinuity. Registration of sharp lightness discontinuity generally causes perception of surface property discontinuity. Hence lightness constancy reveals how the perceptual system operates on mechanisms that actually reflect statistical facts about natural environments — it illustrates how the perceptual system is generally capable of forming accurate states.

According to Burge, the process of forming a state given a luminance contour does not rest on a computation of the probable cause of the contour: Rather, a given contour will yield a determinate state for every case of registration and formation. Formation laws reflect facts about the environment. They are examples of '(…) environmental patterns that that have been *encoded* by the patterns of psychological transformations (…)' (Burge 2010:346, italics mine). Precisely how environmental patterns *are* encoded without being represented is not obvious: However, Burge presents a number of considerations regarding how we should not think of the formation principles: The laws are computable, but they are not computations. The perceptual system does not computationally infer on the basis of statistical facts about the environment. Rather, the facts are reflected in law-like patterns of formation. They are not represented in the system (Burge 2010:346). Law-like patterns of formation can be described for the perceptual system. But *principles of formation* are not accessed by the system (Burge 2010:346).

Another example Burge provides in order to explain how the perceptual system reflects statistical facts about the environment is his account of planar slant/planar surface texture.

Planar Slant from Planar Surface Texture

A sheet directly in front of, and at an angle perpendicular to the visual field of an observer, with a pattern of evenly distributed circles of the same shape and size will appear upright to the observer. If the sheet is slanted backwards, the top circles in the pattern will create retinal images of circles that are increasingly elliptical. The angle will also affect the size of the projected image of the pattern (Burge 2010:359).

The slanted sheet will appear to have the same pattern as it did when it was perpendicular to the observer: Some part of the perceptual system *reflects* environmental facts about the projection of images when forming perceptual states indicating slant (Burge 2010:359).

A sheet that is directly in front of and perpendicular to an observer, with a systematically uneven distribution of non-uniform ellipses will appear to be slanted. Determining the slant of a surface rests to a large extent on registering facts about the texture of the surface. If a surface has some pattern, the structure of this pattern will affect the appearance of slant. Statistical environmental facts (which patterns are more likely to come about in nature, and how certain patterns will reflect light and thereby project images on the retina when slanted) is reflected in the perceptual systems indication of slant. In nature there is a statistical likelihood in favour of an object having a surface texture that creates a pattern of evenly distributed objects of roughly the same size, over it having a surface texture creating a pattern of a systematically uneven array of non-uniform shapes. So there is a statistical likelihood that images on the retina corresponding to descriptions of the latter kind are images

produced by slanted surfaces. The perceptual system automatically represents as of a slanted surface (statistically appropriate cause) rather than upright surface with unusual pattern (statistically inappropriate cause). The perceptual system mirrors natural elegance and order: Textures that create homogenous patterns are statistically more common. The system reflects this principle.

Burge's accounts of lightness constancy and planar slant/planar surface texture accord well with Bayesian models. They illustrate how the perceptual system reflects statistical facts about the environment. However, as the system has been shaped by evolution it reflects *utility*. Perceptual processes indicate as of likely causes of proximal stimulus. However, they do this only in so far as it is '(…) beneficial to animal function' (Burge 2010:345).

The statistical regularities the system reflects are not *accessible* to the perceptual system (Burge 2010:97, 346). They are not applied in any implicit (or explicit) operation of reasoning within the perceptual system (Burge 2010:97). Law-like transitions within the perceptual system can be described mathematically in the same way that any law governing any process can. They are computable, in the sense that they can be given a computational account (Burge 2010:94-95). But they are not themselves computations actually occurring within the perceptual system.

Burge's account suggests that perceptual science does not *rely* on analysing perceptual processes as computations or inference. The way he presents *lightness constancy* and *planar slant,* perceptual constancy does not rely on a *computational theory* of the perceptual system. There is no determinate evidence in favour of the view that perceptual processes represent information about the environment and performs probabilistic

computations on it. Bayesian approaches illustrate this very point. They do not commit to any specific interpretation of the *nature* of actual perceptual processes.

Objective representation of the physical world is the primary kind of representation. Hence, it is evident that perceptual processes cannot be inferential or computational, according to Burge. He expands on this idea when he presents a family of view's that have required inappropriate representational capacities from individuals.

3.2 Individual Representationalism

Theories assuming that objective representation of the physical world depends on representation of conditions such representation exhibit what Burge labels *Compensatory Individual Representationalism* (CIR) (Burge 2010:111) These view's fail to acknowledge that objective representation of physical subject matter does not depend on prior representation. Examples of inappropriate conditions for objective representation are the acquisition of certain language skills; or the ability to distinguish reality from mere appearance; or self-consciousness (Burge 2010:19). Individuals would not represent subject matter *beyond their idiosyncratic perspectives, without such primary* representational capacities, according to certain CIR-views. Inferentialism as Rock presents it, illustrates this: According to his view, perceptual systems must represent information about the physical environment. They would have to *represent* conditions, or *rules of interpretation,* in order to represent the physical world. Hence, perception would not be the initial kind of representation there is. CIR-views have generally placed inappropriate restrictions on which individuals should be thought of as perceivers. Some theories exclude creatures without conceptual thought and language. Some exclude all non-human creatures as well as human

infants. Burge argues that it is empirically proven that all mammals, and certain non-mammalian creatures have perception (Burge 2010:102). There is empirical evidence suggesting that creatures do not need language, concepts or self-consciousness in order to have objective representational states.

3.3 The Principles of Anti-Individualism

Anti-individualism figures in Burge's rejection of CIR-theories. It is a view about the constitutive conditions for an individual to be in perceptual states (Burge 2010:61). Some CIR-views are anti-individualistic. However, anti-individualism plays a fundamental role in establishing an alternative conception of objective representation of the physical world as the primary kind of representation.

Anti-individualism has consequences for what perceptual states are. It represents a philosophical approach to perception: It is a theoretical standpoint regarding the nature of perceptual processes. Burge's aim is to show that perceptual psychology, which does investigate actual perceptual processes, must assume this philosophical stance. Anti-individualism about perception is, according to Burge, not merely compatible with perceptual psychology: It is a working hypothesis of the psychologists (Burge 2010:98). Anti-individualism has two central features: It suggests that perceptual processes reflect environmental principles. It alto supposes that perceptual states have representational content.

Principles (A) and (A') present a general formulation of anti-individualism.

(A) The natures of many mental states constitutively depend on relations between a subject matter beyond the individual and the individual that has the mental states, where relevant relations help determine specific natures of those states. (Burge 2010:61)

(A') The natures of mental states that empirically represent the physical environment depend constitutively on relations between specific aspects of the environment and the individual, including causal relations, which are not in themselves representational; the relevant environment–individual relations help determine specific natures of the states. (Burge 2010:61)

Mental states that empirically represent the environment are any kind of mental state that is about the physical world. (A') claims that individuals' non-representational relations to their environments determine specific natures of such states. Even higher cognitive states *that rely on perceptual representation* again rely on non-representational relations between individual and world. After all, these non-representational relations are necessary for perception. Processes that are not representational account for the processes that cause objective representation. This is why perception is the beginning of representation. Causal relations between the individual and her environment account for objectivity. No additional representation for conditions of objectification is required. As mentioned, as and as evident from (A') causal relations determine perceptual contents.

According to Burge, perceptual systems reflect environmental facts because they have *developed while* creatures have interacted with their environments (Burge 2010:70, 320,326). Objectivity relies on individuals' causal interaction with specific environments. Regularities, principles within those environments are mirrored in perceptual processes.

Burge explains how causal interaction determines objective representation by appealing to a notion of *whole animal function.* Perceptual systems have developed so as to facilitate agency for individuals. Individuals generally need to relate to the *actual*

environment in order to fulfil *whole animal functions.* (Burge 2010:320) Whole animal function is something like the organism's self-sustainment and fulfilment. It is the organism's striving to survive and reproduce. In order to fulfil *whole animal function,* individuals must fulfil biological needs such as eating, navigating and mating (Burge 2010:320). Aspects of the real environment of an individual – how things really *are* in that environment will naturally have an effect on how individuals can fulfil their biological needs. Hence, as individuals interact, non-perceptually, with the environment in order to eat, navigate and reproduce, their 'agency' will be influenced by that environment. The environment will:

> (…) figure in (…) individual's responses to the environment in fulfilling basic needs and activities. (Burge 2010:321)

As perceptual systems develop in accordance with animals fulfilling basic biological needs, aspects of the environment that play significant roles in the fulfilment of these needs will be reflected in perceptual systems. Perceptual systems function to facilitate *whole animal function.* Hence, their processes reflect the actual environment in which animals have eaten, navigated and mated. Actual environmental entities *figure in* perceptual states, in this way. They determine the processes that cause the states. This is evident in Burge's accounts of lightness constancy and planar slant/planar surface texture. Regularities in the perceptual system reflect regularities in the environment.

3.4 Representation

The idea that perception has developed within a specific environment does not in it self suggest that perceptual states have representational content. It merely suggests that real entities, with biological relevance will affect the structure of perceptual processes. Burge argues that

scientific explanation of these processes makes *non-trivial* appeal to representational *contents* (Burge 2010:292). Perceptual science assumes that its' explanada are states with veridicality conditions. The science appeals to representational content as a 'real kind' (Burge 2010:293).

Perceptual science makes *non-trivial* appeal to representational content in the sense that these explanations could not have been replaced by non-representational notions (Burge 2010:293). Hence the central argument in favour of representationalism is a kind of non-reductivism. The notion of a *representational function* figures in this argument. It reveals how Burge describes perception as a *teleological* process.

3.5 Teleology

A principle that is closely related to anti-individualism illustrates the notion of a representational function. This principle does not *follow from* anti-individualism (Burge 2010:68), but it is closely related.

(B) For an individual to have any representational state (such as a belief or perception) as of a subject matter, that state must be associated with some veridical representational states that bear referential, indicational, and attributional representational relations to a suitably related subject matter. (Burge 2010: 68)

Burge claims that all representational states constitutively depend on veridical representation. Perception *functions to* represent. Understanding what a functional capacity is, relies on understanding '*successful realization'* (Burge 2010: 68) of that functional capacity. The idea is that representational states *are the states they are* in virtue of standing in some relation to states that have fulfilled their function. Understanding what perception *is* relies on understanding instances upon which the perceptual system performs the task it functions to, *namely to veridically represent.*

Hearts, hammers and ovens are 'functional systems'. These entities have specific tasks they function to perform. Understanding a heart, or a hammer, or an oven requires understanding what these things *function* to do. Indeed, being a heart, hammer or oven, seems to presuppose that the heart, hammer or oven *can, sometime has, or is appropriately related to something that has,* performed their appropriate function. Being a heart requires that the entity in question is either in relation to other hearts (in the sense of being the result of a developmental process) or that it is constructed in order to perform the functions of a heart (take artificial hearts — artificial hearts are also appropriately related to things that actually have performed the appropriate tasks that hearts perform.)

Representational function is 'constitutively associated with representational success' (Burge 2010:309). Representational contents are the contents that they are in virtue of their accuracy conditions: In virtue of what they *aim at representing*. We *individuate representational contents* in virtue of what they aim to represent. For example, sentence 'George is happy' means (represents) what it does in virtue of its conditions for being true. The sentence *functions* to represent that content.

Principle B claims that perceptual contents are the contents they are in virtue of being 'related to veridical states that bear (…) representational relations to a suitably related subject matter' (Burge 2010:68). Hence perceptual states are the representations they *are, they have the content they do,* in virtue of being related to subject matter. Their representational functions depend on their relations to the things they aim to represent: They have their content in virtue of relations to those things.

As perceptual states are *type individuated* through their content, they have representational functions. Only representations have such representational functions. Non-psychological, biological systems (such as the immune system) do not *aim at accuracy* (Burge 2010:300-301). Biological systems have functions. However, they do not function to be accurate. They aim at *success,* at 'functioning well enough for to contribute to survival and reproduction' (Burge 2010:303). No non-mental system aims at accuracy. That is why representational contents are not obviously reducible to non-representational phenomena. Non-representational phenomena cannot capture the teleological aspect of representations. They cannot capture the *specific teleology* that only representations have.

Notable theories of intentionality have attempted to explain representation in terms of some notion of biological function. Teleological theories of mental content generally attempt to reduce mental representation to other kinds of *functional* capacity (Burge 2010:299). Notably, such theories[8] aim at explaining the representational properties of mental states in terms of biological function. A simplified explanation of such a theory might suggest that what it means for a mental state to be about a cat *is* that it was selected for in order to be about that cat. On this very simplified explanation, function is supposed to *explain* representation. According to Burge, teleological theories aim to explain what objective representation *is* in terms of some notion of '(…) co-variation or causal co-variation, or structurally isomorphic causal co-variation' along with 'biological function' (Burge 2010:194). Hence these theories aim to *reduce* representational notions used in psychology to causal correlation between a representation and the things it represents. Psychological explanations would not *rely* on a notion of *veridicality* according to such

[8] See Neander (2012) for a good introduction to teleological theories.

a view: Veridicality, accuracy, representation can be explained in terms of other, not exclusively psychological, notions. A representational state's being *accurate* about some subject matter can be assimilated to how a heart can succeed to fulfil *it's* biological function. Being true is just being *selected for correlation,* on this view. Burge's account does not aim to reduce representation.

Perceptual systems do *have* biological function, according to Burge. The way I read Burge, the reason it that the *function* of a representation *is* to be veridical. Hence, explaining what the *function* of a representation is relies on understating *successful realization* (Burge 2010:68). Representational states do not merely aim at correlation, according to Burge. They aim at *accuracy.* Hence, understanding the specific function that a representation has relies on some notion of *veridicality.* A representational content constitutively relies veridical representation (Burge 2010:68). Perceptual processes have been selected for so as to cause states that function *to be accurate about* some subject matter. Not merely to *correlate* with that subject matter. This *specific* functional aspect that only representational states have, cannot be explained in terms of *function* and *correlation* alone.

Reflecting on the notion of a function, there is a sense in which this point is quite intuitive. How could a functional capacity be *reduced to something else?* Can the specific function of a hammer, a heart or an oven be *wholly* explained in terms of anything else? I think there is an intuitive notion of a function where *no* function can wholly explained in terms of something else. However, there is a sense in which a hearts, or a hammers or an ovens function *can be* reduced: We can describe what a specific hammer, or a heart or an oven does when it is executing its function in non-functional terms. We can give explanations in terms of for example physics – explanations that make no appeal to the specific

function of a hammer, or heart or oven in order to explain what a specific hammer or heart or oven is doing in each instance of its performing its function. The way I read Burge it is this kind of reductive explanation that *cannot* be given for a representational content. If the function of a representational state is to aim to be about something, it is not obvious that it's what it does when it executes that function can be described without making use of some kind of notion of veridicality. Explaining what a representational state is doing depends on describing that specific representational function — namely aiming for accuracy. The specific way that a representational content functions to be accurate cannot be captured without making use of some notion of veridicality, according to Burge. Hence he argues in favour of non-reductivism about representation.

3.6 Evolution

Where does the representational function of perceptual states come from? If biological systems aim at 'contribution to fitness' (Burge 2010:303) alone, what makes a perceptual system *aim at accuracy?* After all, these systems have evolved. Why would systems aimed at accurate representation evolve from non-representational structures? Why would the law-like processes that govern perceptual formation reflect actual environmental statistical regularities? Couldn't perception simply be *wholly interest-dependent?* Natural selection has shaped the perceptual system, as evident from how these systems reflect nature. Why would this selection aim at accuracy?

> Evolution does no care about veridicality. It does not select for veridicality *per se.* Being fitted to successful evolution is a matter of functioning well enough to contribute to survival and reproduction. Well enough often coincides with veridicality. (Burge 2010:3030)

The quote expresses what maximal fitness observers (Geisler & Diehl 2003:396) reveal[9]. Accurate perception is beneficial for survival. Evolution itself in no way recognizes *veridicality.* However, individuals with accurate perception will in general be well equipped to sustain whole animal function. Navigating, eating and mating rely on interacting with environmental entities. Successfully navigating, mating and eating is in general facilitated by *accurate* indication of actual environmental entities. As mentioned, interactions between individuals an environment, *individuals' striving to fulfil whole animal function,* are the causal relations that determine perceptual representation. In this sense, Burge's account does appear to entail that perception is '*wholly interest dependent'.* However, interest and accuracy will generally co-vary. Hence, the fact that perception is aimed at contributing to successful behaviour suggests that it is aimed at objectivity.

> The notions perceptual state and perceptual system are partly teleological notions. (Burge 2010:309)

Perceptual systems are teleological because they function to represent. Like any functional part of natural systems, they have been selected for, in order to *fulfil that function*. It is useful for an organism to have a perceptual system that generally fulfils its function, just as it is useful for an organism to have a heart that generally fulfils it's natural function.

As certain phenomena (such as rabbits' 'thumping') reveal, veridicality and utility do not always coincide. However, reflecting on what agency *is,* simply suggests that it generally will.

[9] See p. 26. for discussion of maximal fitness observers.

Action is prior to perception (Burge 2010:326). Not all kinds of action depend on intellectual capacities, such as desire, intention or will. Organism's striving toward fulfilling whole animal function involves a kind of agency. This kind of agency has determined the shape of perceptual processes. These processes determine which specific physical particulars; entities and subject matter are represented in perception. This is how perception is anti-*individualistic*. The agencies of our ancestors determine what we perceive.

Burge's account of perception as selected for a specific kind of teleology (its aim at accuracy) presents a way of giving answer to the traditional subjectivity/objectivity problem I mentioned in chapter one. How do perceivers, from their subjective *perspectives,* perceive the *real world?* The answer is that evolution has selected for systems that generally do. These systems have developed so as to reveal how things *really are*. The processes that govern perceptual formation are molded by nature to *reflect reality*. Hence, perceptual states represent a mind-independent world. Biological needs have determined the shape of perceptual processes. A capacity to relate to how things really are beyond an individual will generally promote her ability to fulfil such needs.

3.7 Perception as Experience

Mental representation is a phenomenon that philosophers have aimed at reducing to non-representational phenomena. This reflects the general desire in philosophy to reduce mental phenomena to non-mental phenomena, in order to understand how the mind fits into the natural order of things. Burge's project presents a non-reductive way of 'naturalizing' representation. He claims to have found mental representation in natural science. Hence, he gives a non-reductive explanation of intentionality which accords with natural science.

Intentionality begins in perception (Burge 2010:10, 3-108). It is *shaped* by the nature around us.

Representational states are perspectival. They contain a certain mode of representation, or *way of referring*. According to Burge's picture, perceptual representation does not constitutively rely on consciousness (Burge 2010: 368, 374). Perceptual states can be (and sometimes are) *unconscious perspectives* on this view. Such states unconsciously represent the world to subjects in specific *modes,* or ways of referring. The idea that perceptual processes and have developed so as to sustain *whole animal function* plays an important role in Burge's explanation of how perceptual *perspectives* are not necessarily experiences. In the following chapter I discuss a view of mental representation according to which Burge's view overlooks a part of the nature of mental representation.

Ch. 4 Experience and Intentionality

This chapter is a discussion of the relation between consciousness and mental representation — in perception and in general. I explore a strong thesis of phenomenal intentionality. According to this view, only conscious experiences are mental representations. I assume that we have unconscious mental lives. What I wish to explore is a view upon which this unconscious mental life contains no representational content with *veridicality conditions.* According to the view I present, neither consciousness nor mental representation is the mark of the mental. Some non-representational and non-experiential property must delineate the beginning of mind.

I do not make in-depth discussion about what this property should be. I think that a promising suggestion for what such a property could be, is a kind of intentionality, or representation, that does not rely on veridicality-conditions (i.e. that does not rely on *way of referring* or *mode of reference*). I briefly discuss this in 4.7.

As mentioned, in this thesis I have simply assumed that intentional states have representational content (that they have conditions for being accurate). I will continue to do so in this chapter. Hence, the strong thesis of phenomenal intentionality is only meant to address sates with accuracy conditions.

According to the strong phenomenal intentionality I explore, only *experiences* manifest this kind of representation. Hence, I explore a view of intentional states as conscious. It does not entail that all experiences are mental representations. Only that all mental representations are experiences.

I do not argue in favour of this strong view of phenomenal intentionality. Rather, I explore it as a coherent thesis. I present my primary motivation for being attracted to it: Namely that the view might capture how mental representations are *perspectives* in a sense that opposing views do not. Mental representations are *perspectival,* in the sense that they are specific modes of representation: They are specific ways of referring (Burge 2010:37). I suggest that all mental representations are *perspectives* in the sense that they are subjective experiences. I suggest that what makes a mental representation an individual's perspective is not merely that it has representational properties. Rather, it constitutes an individual's perspective in virtue of being her *subjective outlook* in a way that only experiences are. It manifests itself in an individual's psychology as an experience. According to the view I present being *subjectively manifest* in this way is a necessary condition for being a mental representation. It is not the case that states become perspectives because they have *perspectival representational content.* Rather, they have perspectival representational content in virtue of being the perspective *of an individual.* Being the perspective of an individual is, I suggest, a sufficient (but not necessary) condition for being an experience.

Reflecting on the notion of a *mode of representation, or way of representing* is the central ambition of this chapter. Smoke represents fire. Leaves represent autumn. However, only mental states with representational contents represent subject matter *in a specific way, from a point of view.* The strong phenomenal intentionality view is first and foremost a view about the nature of *veridicality.* It simply is the view that only experiences represent from a point of view. Hence only experiences can have specific conditions for being *veridical* or *accurate*[10].

[10] The content of a sentence has truth-conditions. What I take for granted is that the content of a sentence is mind-dependent — Only when understood by a conscious perceiver does a sentence *mean* anything. Hence the content of a sentence, when understood, exists within the psychology of an individual as a mental representation. In this sense, only mental states have conditions for accuracy. There is no third realm where contents exist. Any content is *part of someone's psychology.* The strong phenomenal intentionality can be understood independently of this point. I does not hinge on the

In (2010) Tyler Burge explains a way of thinking where perceptual states are the perspectives of individuals (i.e. are mental representations that should be attributed to the individual's psychology), not in virtue of being experiences (Burge 2010:376). Experience does not account for how representational states 'belong to' or 'represent for' the individual. Rather, the relation between agency and perceptual processes (i.e. how perceptual states have developed in order to sustain *whole animal* function), and the nature of what it is to be a perspective explains why we should attribute perceptual states to an individual, and not merely to her perceptual system. Since perceptual states are *perspectival* states of the individual not in virtue of being experiences, this account is compatible with the view that there can be *unconscious perspectival states of the individual.* Burge argues that there is reason to think that there are unconscious perceptual representations (Burge 2010:375).

While I agree with Burge that perspectival states belong to the individual not *in virtue of* being experiences, I suggest that the only way in which perspectival states belong to individuals is *as their conscious perspectives.* Hence, while the role perceptual representations play in explanations of agency and the causal-evolutionary background of perceptual processes explain *why* perceptual representations are the perspectives of individuals, I suggest that an explanation that makes appeal to these features alone is silent on *what* a perspective of an individual *is.* I suggest a way of thinking where being the perspective of an individual simply entails being an experience. Any perspective is an experience, on this view. While one can explain how perspectives come to belong to individuals without appealing to conscious experience, any such perspective *is* an experience according to strong phenomenal intentionality.

---

idea that all content is mental in this sense. In it self the view only entails that all mental representations are experiences.

The view relates to the discussions in the previous chapters of this thesis because it suggests that, while perceptual psychology makes use of intentional notions, the psychology has not discovered that there is any unconscious representation. Perceptual psychology has, the view suggests, only discovered that perceptual experiences are mental representations.

The chapter is intended to present and explore the foundations for a view of the mind. The view I wish to explore is that representational mind cannot be separated from consciousness: Mental representation is always subjective, in a sense that only experiences are. It relates to contemporary theories of embodiment. Unfortunately it is beyond the scope of this thesis to address such theories.

I begin by further addressing what I take 'experience' to be. I go on to present first a weaker form of phenomenal intentionality than the one I wish to explore. Subsequently I explain the strong version.

4.1 What is Experience?

Experiences are not merely events that a subject undergoes in this context. An experience is something that has some character that is presented for, or to, the subject of the experience. An experience is constitutively subjective. It has subjective character: It is like something for the experiencing individual[11]. Experiences are transparent to the individual. They are, in a sense, directly present to her. This does not, I think, entail that an individual is required to attend to or be able to conceptualize about her experience. However, experience is constitutively *subjective* mental event, in the sense that it is continuously like something qualitatively, for the subject.

[11] See Nagel (1986)

Individuals have consciousness. They are conscious only when they experience. Hence, consciousness is first and foremost something that *individuals have* when they experience.

A specific kind of subjectivity is what makes something conscious, on my view: Having a *subjective experience* just means being conscious. Subjective character[12] can be thought about as *what separates being awake from a dreamless sleep* or *whatever enters into the stream of consciousness.* There is, on my view, no distinction between the subjective character of an experience and the phenomenal character of the experience. Something's seeming blue (in a phenomenal sense) constitutively depends on it's seeming blue to a subject. Only in someone's experience does something seem blue. I see no reason to differentiate between 'conscious event', 'experience', 'conscious experience', 'phenomenally conscious event' or 'phenomenally conscious experience'. Phenomenal properties are properties of experiences. Subjective character and phenomenal character are identical. They are properties of experience. The notion of experience I am delineating does not recognize a difference between phenomenal consciousness and access consciousness[13].

4.2 Phenomenal Intentionality

Views that advocate phenomenal intentionality, suggest that all intentionality somehow relies on or derives from phenomenal properties. Without experience, there would be no intentionality. Phenomenal intentionality as a general thesis is given thorough explanation in Kriegel (2013). He suggest that the thesis is based on:

[12] See Kriegel (2005)
[13] See Block (2003:483)

> (...) the idea that intentionality is injected into the world with the appearance of a certain kind of phenomenal character. It is when the relevant phenomenal character shows up that intentionality makes its first appearance on the scene. (...) once this phenomenal character appears (...) intentionality can be 'passed around' to things lacking this (or any) phenomenal character. But the source of all intentionality is the relevant phenomenal character. (Kriegel 2013:3)

While Kriegel does not presuppose that intentionality is representation when they provide this definition, I will continue to assume that it is. According to Kriegel theories that fall under the phenomenal intentionality research program propose that states get their intentional properties from some kind of phenomenal character. His explanation is silent on just what this phenomenal character is. The most general thesis of phenomenal intentionality simply suggests that all intentional states are intentional in virtue of standing in some appropriate relation to some conscious experience. My ambition is to explore a stronger kind of phenomenal intentionality than captured by the quote. This stronger view entails that intentional properties appear in experience: Only within someone's experience does proper intentionality exist.

The view is connected to what Kriegel labels *intrinsic subjectivity* (Kriegel 2013:11). Intrinsic subjectivity has been put forward as a property that separates phenomenal intentionality from other kinds of intentionality.

The idea behind the notion of intrinsic subjectivity is that mental representations are intrinsically subjective only if they *intrinsically* represent *to* or *for* someone (Kriegel 20103:11). Representation can be framed as both a two- and three-place relation. Spots can represent measles; '4' can represent 4; a picture can represent a person. These are all ways of talking about representation as a two-place relation. However, spots can represent measles to the doctor; '4' can represent 4 to a math-student; a picture can represent a person to an audience. These examples illustrate how representation is a

three-place relation. States with intrinsic subjectivity, instantiate a three-place representational relation of this kind, *in themselves*. These states themselves are constitutively three-place relations. The representational states in question contain in them an element of being representations *to,* or *for,* an individual (Kriegel 2013:11). They states *intrinsically* represent for someone. They are not representational states that can be interpreted by someone in another state: The very representational state in question essentially contains within it that it represents to a subject. The suggestion that *only* phenomenally intentional states have intrinsic subjectivity is prevalent in theories of phenomenal intentionality. The idea is that only experiences intrinsically represent to an individual.

The claim that only phenomenally intentional states have intrinsic subjectivity relies on the idea that experiences have subjectivity of special kind. This kind of subjectivity must be *in place* if a state is supposed to intrinsically represent to someone.

The strong view of phenomenal intentionality relates to the thesis of intrinsic subjectivity in this way: It entails that *all* mental representation has intrinsic subjectivity of this kind. For anything to be a real mental representation, it cannot only represent a subject matter — that state itself must represent subject matter *to a subject.* Representing to a subject is, according to the view I describe, equivalent to being an experience. Only experiences, I suggest, are subjective in the appropriate sense. Whenever I make use of the term *intrinsic subjectivity* I assume that this is an experiential phenomenon.

The reason why I am drawn to the idea that all mental representations are *intrinsically subjective* is that intrinsic subjectivity can capture what I think is an essential feature of a *perspective.* It captures *how* perspectives belong to subjects. Perspectival states are not only states that can be attributed *to an individual.* Arguably, any number of non-perspectival states can be. For

example, *behavioural states such as* walking or sitting down or drinking are states that are attributable to *me*. I do the walking, sitting or drinking. However, these states do not constitute my perspective. Perspectival states, such as a visual experience of a snowy field, or my thoughts about my future, appear to belong to me in some way that states of walking do not. So what makes something a perspective?

My suggestion is that that all genuine perspectives manifest these two essential features: 1. They are representational contents. 2. They are subjective experiences. All perspectives instantiate these two features together, according to the view I am presenting — only perspectives instantiate the first one. A perspectival painting is not a perspective other than metaphorically so: However, someone's visual experience of the painting is. A speech written from 'the point of view' of Molly Bloom is not a perspective other than metaphorically so. However, someone who reads Molly's speech can have Molly's metaphorical 'perspective' as their real perspective when they read it.

4.3 Perspective as Representation

Representational contents are perspectival. They constitute *modes of presentation, ways of representing*. This perspectival feature of representations is, I take it, what Burge means by *mode of reference,* or *way of referring* (Burge 2010:37-38). I think Fregean senses capture this feature (Frege 1892:217-218). Representational contents provide a specific kind of 'outlook' or 'way in which' something is represented. No accuracy conditional representation of anything is *complete*: There is no representational view from nowhere. Mental representations are constitutively perspectival. They present some subject matter in a specific way. Perceptual representations of physical reality, for example, represent entities from specific angles in specific sensory modes. Perhaps thought, belief and desire are modes of a similar kind to the

sensory modes. Sensory mode or cognitive mode does not wholly capture what a mode of presentation is for mental representation. Not only are visual representations different from auditive representations, not only are representational thoughts different from representational desires. Thoughts, desires, beliefs, perceptions: Such mental representations also refer to specific parts of subject matter in specific ways. They pick out specific 'chunks' of reality in determinate 'fashions.'

Representational mental states are perspectival in the sense that they refer to, or represent subject matter *in a specific way.* Being a mode of presentation, or a *way of referring* is the first feature I suggested as a feature of all and only perspectives.

The obvious question at hand is: What more is there? What more is there to perspective, than being a mode of representation? The idea I wish to express is essentially this: In order to *be* a mode of representation, something has to be a mode of representation *to* or *for* someone. A constitutive part of being a mode of representation is that it somehow belongs to a subject. A perspective is something that someone has. My idea is that the best way of capturing this, indeed, the only way that doesn't loose or overlook something, is by suggesting that any mode of representation is subjective in the sense that is has a subjective character. This is why I suggest that all mental representations are intrinsically subjective: They are all perspectives. My suggestion is that perspectives are the perspectives of individuals in the sense that they are essentially their outlooks or points of view, and that these outlooks or points of view constitutively have some way of being for the subject of the perspective.

4.4 Manifestation and Subjective Character

Strong phenomenal intentionality of the kind I explore is a view about constitutive conditions for modes of presentation. The key claim is that any *mode of representation* must manifest itself in an individual's psychology *as an experience.* Modes of representation cannot *exist in the psychology of the individual* if not as an experience. They have to be '*embodied'* as experiences. Experience is in a sense the structure that *realizes* representation. These considerations are opaque. However, I think they do reveal something about the very nature of a mode of presentation. Modes of presentation are *angular.* They provide specific, determinate ways of picking out subject matter. That states in the psychology pick out specific subject matter seems to me to somehow *rely on* that subject matter having some kind of internal *qualitative way that it is,* in the individual's psychology.

What is the subjective character of mental representation? In what *manner* do representations manifest or embody? For *perception* I think a plausible suggestion is sensory feel. By sensory feel, I mean the qualitative character that characterizes ordinary perceptual experience. I mean the kind of bodily sensation that is associated with hearing, seeing, touching etc. Sensation is prior to perception. Sensory information registration precedes perceptual representation. Descriptions of sensory information registration in biology and perceptual psychology do not make any essential appeal to sensory feel.[14] Supposedly, sense perception can be unconscious. Sensory feel is conscious, according to the notion of consciousness/experience that I stipulated in the beginning of the chapter. Sensory feel 'has some character that is presented for, or to, the subject of the experience,' and is 'like something, for the experiencing subject.'[15] There is nothing *more* to a sensory feel, than its' subjective character. At present there seems to prevail something like a kind

[14] For example, the explanation of retinal information registration in Wade & Swanston (1991) does not mention anything like *sensory feel.* See p. 4 – 6.
[15] See p. 48 where I expand on experience.

of agnosticism *in science* regarding whether sensory information registration and perception has sensory feel. We just do not know whether fish *feel* pain[16], or whether bees[17], even if they should count as perceivers capable of objective representation of physical reality, have any subjective experience. The nature of subjective experience *is* its subjectivity. We can only *assume* that others experience. Hence, agnosticism regarding sensory feel does not, I think, suggest that sensation is not subjective experience. I am interested in the subjective character of *mental representation.* Hence, I will not argue that all sensory information registration has sensory feel to it. What I will suggest is that the way a perceptual representation manifests itself in an individual's psychology is *as a subjective experience.* (I do not commit to the view that all perception is experience. Only that perceptual representation is.)

As I have mentioned, I do not think that subjects need to attend to or be able to conceptualize about their experiences. Hence, my view does not entail that an individual must be able to entertain any kind of meta-perspective on her experience in order to have the experience. All that is required for something to be an experience is that it has some subjective character. My suggestion is that all perceptual representations manifest themselves in an individual's psychology as experienced sensory feels. These feels have *subjective character.* If state has sensory feel it fulfils a sufficient condition for being an experience. Again, my suggestion is that any mode of presentation must manifest it self in the psychology of an individual in some kind of experiential *manner,* in order to be her perspective. For perceptual modes of representation I think such modes are sensory feels. Hence, sensory feel constitutes the *intrinsic subjectivity* of perceptual representations. They *allow* a state to represent *to* or *for* an individual. They *embody* perceptual representations. In a sense they are the 'matter' that perceptual

[16] The reason *why* we do not know this is, partly, that we just do not know what a sensory feel or conscious experience *is!*

[17] Burge argues that we simply do not know whether certain perceivers, such as bees and spiders are conscious. (Burge 2010: 375)

representational contents are *made of.* They are something like the *symbols* that carry a perceptual states representational content. They explain *how* a representational content exist in the subject's psychology as the perspective of a subject.

I think non-perceptual representations have subjective character. I remain agnostic as to what the subjective character of non-perceptual states such as beliefs, thoughts and desires is. I am inclined toward thinking that such states have a kind of sensory feel. I think there is good reason to investigate the idea that the subjective character of perceptual experience is (similar to) the subjective character of conscious belief, thought and desire. On such a view, any mental state has a certain *bodily* character to it. This is closely linked to contemporary theories of embodied cognition, however discussion of the idea goes beyond the scope of this thesis.

As mentioned, the strong phenomenal intentionality I have presented in no way commits to the idea that there is no unconscious perception, belief, thought or desire. It only entails that such unconscious states are non-representational. At least, they cannot represent in the sense of having representational contents, or modes of representation. Modes of representation are constitutively perspectival. Perspectives are manifested in the psychology of an individual as a state with subjective character, according to the view. As mentioned I briefly discuss unconscious mind in 4.6.

The notion of experience that underlies this discussion is perhaps not a very strong one. Anything can be an experience, on this view, if it has some subjective character. I think it is possible that an individual is not capable of expressing the subjective character of her experience. I think it is possible that she might not attend to or even notice the subjective character at times. For example, an individual might not continuously notice the flavour of her gum. This does not, I think, entail that the flavour of her gum cannot be part

of her 'waking life' or 'stream of consciousness' in these cases. I think it only entails that she is not always cognitively attended to or in contact with her experience. I think that there are cases where her gum still has a flavour to her. Subjective character constitutes experience. Subjective character needs not be cognitively available. But it is constitutively always like something for a subject. Hence, in a sense, the reason why I equate experience and representational content is that I have doubts concerning whether representational contents can *fail to meet the requirements* for being an experience. A representational perceptual state is a *mode of presentation* that is constitutively sensory. It is a *perspective* given by sensory information, with specific sensory 'make up'. Perceptual representations are presented by one or several sensory modes. They belong to the individual, as her sensory *perspective.* Can a sensory perspective really fail to be *subjective,* in the sense that it has characteristic subjective *feel* to an individual? What more is *required* for having subjective character? What *gives* subjective character, if sensory perspectives do not all have it? In a sense, the reason for thinking that perspectives are experiences is two-fold. Not only does the perspectival feature of a mode of presentation suggest that representations are *subjective.* Reflection on experience seems to suggest that perceptual representational states *are* experiences because they meet the *requirements* for being such.

In (2010) Burge presents a way of thinking about how perceptual states are the perspectives of individuals (Burge 2010:371), that makes no appeal to experience. While I sympathize with his general view, I believe it fails to capture *how* perceptual states *manifest* within an individual's psychology as her perspective. In the following section I expand on this view.

4.5 The Subsystem/Individual Distinction

As I explain in chapter three, perceptual states are the results of processes that have developed in accordance with *agency.* According to Burge, agency is

distinct from things that merely happen *within,* or *to,* an individual (Burge 2010:327). Action and behaviour lie at the level of individual, not merely at the level of internal sub-systems. Digesting is not an action. Ingesting is. (Burge 2010:328) As primitive agency has shaped perception, perception *belongs to the individual.* Perceptual states should be attributed to individuals, and not *merely* their perceptual systems, according to Burge (Burge 2010:374). The relation between agency and perception in part explains the representational nature of perceptual states. Perception has developed in order to aid *individuals* so as to facilitate their actions (ingesting food, for example). This is why they represent the things they do.[18] Perceptual states indicate entities that are of relevance for the organism as a whole. They represent such entities *to the individual.* Burge does not appeal to any kind of *subjectivity* in order to explain this. Perceptual representations are not the perspectives of individuals *in virtue of* representing to individuals in experience:

> More importantly, there is considerable evidence that individuals, not merely subsystems, have unconscious perceptual states. So there is reason to doubt that consciousness is constitutive either to the individual/subsystem distinction or to perception. (Burge 2010:374).

Burge provides two reasons why perceptual states are individual level states:

> One resides in the connection between perception and whole animal, or individual, function — paradigmatically individual agency. (...) perceptual kinds constitutively figure in individual functions — in fulfilling needs and guiding actions.' (2010:370)

A second reason is that:

> Being a locus of perceptual representation just *is* being an individual with a representational perspective. The objectification and representational content involved in perception

[18] I expand on this is chapter three, see 3.7.

constitutes a point of view on environmental representata. Representational perspectives or points of view are constitutively attributable to individuals. (Burge 2010:371)

The first reason states that perceptual states are states of the individual because the processes that determine perceptual states have developed in accordance with primitive, pre-perceptual agency. Perception has developed so that individuals might successfully fulfil their biological needs for eating, navigating and mating (Burge 2010:320).

In itself this does not make the states *perspectives* of the individual. Some agential states are not perspectives. The kind of pre-perceptual, primitive agency (2010:320-321, 326-342) Burge describes in order to explain how perception has developed in accordance with agential needs and actions, in no way relies on any representational state. Eating, navigating and mating are examples of agency that does not constitutively depend on perspective.

However, the perceptual system functions to accurately represent for individuals. (Burge 2010:308) This gives perceptual states a special role — they function to indicate entities that are of relevance to individual creatures. Perceptual states are states *of the individual* because they *function* to *aid agency*. Perceptual states are representational. Representational states are perspectival. Hence, perceptual states are perspectives of individuals.

The second reason Burge describes simply states that perceptual states are attributable to individuals *because* they are perspectives. It addresses how perceptual states have the contents they have (are the states they are) because they belong to individuals as perspectives on subject matter.

I think both these reasons express truths about the nature of perception and mental representation. Perspectives belong to individuals. Perceptual states have developed in order to be perspectives. *Perspective* and *individual* are

interrelated notions. Burge's account explains *why,* as in — in virtue of what processes — representational states represent to the individual. He gives two reasons for thinking that representational states *are* states of the individual. However, he does not explain *what it is for a state to be a representational state of the individual.* His view establishes why perceptual states are the perspectives of individuals. However, he does not provide discussion of what a perspective *is.* He does not address the manner in which a perceptual state or any mental representation *manifests itself in the psychology of an individual.*

Burge's account assumes that the only essential property of a mental representation is that it has representational properties and is attributable to an individual. Hence, there can be[19] unconscious mental representations. I think the account appears to overlook the *manner in which* a mental representation exists in the psychology of an individual. Perhaps it presupposes that *being* a representational state is manner enough. I am disinclined to think that this is accurate. My suggestion is that perceptual representations and representations in general exist *as experiences.*

With this in mind let me return to the strong phenomenal intentionality I put forward. According to this view, all mental representations are *intrinsically subjective.* I suggested that only conscious experiences are intrinsically subjective. The strong thesis I put forth simply is the view that all mental representations are *essentially conscious.* Hence, there cannot be any unconscious mental representations. An unconscious state cannot be a representation, on this view. A consequence of the view is that establishing *that* individuals have representational states entails that those individuals have experiences. Claiming that perception is constitutively representational simply entails that perception is constitutively conscious, on this view.

[19] And, according to Burge, there is reason to think that there are, *unconscious perceptual representations.* (Burge 2010:376)

I think Burge's account appears to leave something out in a discussion of perspectival states. His explanation simply reveals why individuals have perspectival perceptual states, without making use of any notion of consciousness or experience. The account is supposed to underlie the idea that perspectival states are not necessarily experiences (Burge 2010:376). However, the way I interpret the view it simply explains which processes give individuals perspectival states. It does not discuss what it *is* to be a perspectival state of an individual. It does not discuss *how* a mental representation can manifest itself in the psychology of an individual. It does not attend to what it is to *be* a mental representation in the psychology of an individual.

I think there is an intuitive sense in which *perspectives* must manifest or embody as subjective characters in the psychology of an individual. For perceptual states, I think it is plausible that this subjective character is sensory feel. Hence, an account that does not discuss the manner in which perceptual representations exist in an individual's psychology is an incomplete account. According to Burge, there are scientific reasons for thinking that perceptual representation can be unconscious.

4.6 Unconscious Representation in Science

Burge provides three examples of scientific results that suggest that there can be unconscious perceptual representational states.

The first reason is that experiments made on humans with certain dissociative disorders, such as blindsight, prosopagnosia and extinction-neglect syndrome, have been taken to suggest that there can be visual perception with no attached phenomenal experience (Burge 2010:375).

Burge also justifies the assumption that there can be unconscious perception by arguing that we do not know whether creatures such as spiders and bee's have experience. However these creatures are perceivers. Therefore, as perception can be '(...) confidently and firmly attributed to bees and spiders without knowing whether they are conscious (...)' (Burge 2010:375) we should not presuppose that they are. Bees and spiders might, for all we know, be examples of creatures with perceptual capacities that lack phenomenal consciousness, according to Burge.

A third example from psychology addresses early stages of vision. These '(...) may count as perception by the individual (...)' (Burge 2010:375), in virtue of manifesting perceptual constancy. Perceptual constancy is objective perceptual representation. However, these early stages of vision are supposedly not necessarily conscious experiences.

The first example addresses phenomena that have been given significant weight in discussions about perception and consciousness[20]. However, there is not scientific agreement as to the significance of these phenomena[21]. It is beyond the scope of this thesis to discuss this disagreement. I will mention that blindsight; prosopagnosia and extinction-neglect syndrome is rare. Individuals do not generally have blindsight. There is reason to question how much import dissociative disorders should be given in discussions of conscious perception.

Burge's second example does not carry much evidential force in favour of unconscious perceptual representation. Why *not* assume that bee's and spiders have conscious experience? What are the reasons in favour of

[20] Other examples of philosophers who make use of blind sight and other dissociative disorders in humans in order to argue that there can be unconscious, representational perception include Jesse Prinz (See 'When Is Perception Conscious') and Ned Block (2007)

[21] See Ian Philips' lecture 'Unconscious Perception Revisited'

reluctance? If a conceptual, constitutive relation between experience and perspective can be established it entails that any perceiver is conscious.

The third example is, I think, the most convincing one. If early stages of vision are representational and unconscious, they simply disprove that representational states are constitutively conscious. However, I think falsifying the view requires much philosophical analysis of the very nature of *conscious experience.* It is difficult to establish whether a mental state *has* subjective character. After all, subjective character is simply unavailable to anyone but the experiencing subject. Hence, increased *philosophical* understanding of the nature of subjective character, perspective, attention, access, bodily feel, perception and though is what primarily informs the plausibility of the view I have presented.

4.7 Non-Representational Mind?

The primary problem for the strong phenomenal intentionality thesis is not, I think, scientific evidence in favour of an alternative view. Rather, strong phenomenal intentionality is faced with having to explain what unconscious beliefs; desires; thoughts and perceptual states are, if not representations. I do not have a full positive account of what unconscious mentality is. Perhaps unconscious attitudes and perceptions are *dispositions* to believe, or think or perceive. Perhaps they are dispositions to behave in a certain manner. Perhaps they are representations without any accuracy conditions — perhaps they represent only in the sense that they function to correlate with certain experiences. Perhaps they are stored information that has been represented in experience. Some kind of *eliminitavism* about unconscious intentionality has to be true, according to the view I present.

I think there are deep connections between subjective character, perspective, representation, veridicality and perception. Hence, I think the strong view of

phenomenal intentionality that I have presented is at least worth exploring further, in relation with science and philosophical progress toward understanding what experience is.

Appendix: Ideological Consequences of the View that Perception is Representational

There are certain consequences of ideological kind that follow from representationalism about perception. It provides some explanatory ground for how minds come to be intentional at all — namely that perceptual processes grant intentional states that provide the background for other intentionality. An important virtue of the view that perception is the beginning of mental representation is that it recognizes the similarities between the mental capacities of humans and other animals. The view Burge advocates, suggests that non-human individuals, as well as infants and adults who lack certain higher cognitive skills, still perceive their environments in ways that are relevantly similar to human adults who possess such abilities. Several CIR-views have, I think, expressed a speciesist tendency to overestimate the significance of capacities such as language and conceptual thought. These capacities are arguably significant in human self-reflection, and they have been extensively used to separate *us* from *the others* — People of other ethnicities, women, humans with cognitive disabilities, children, other animals and nature in general. The same point can be made for many theories that have amplified the significance of conscious experience for perception.

Philosophers and psychologists have stressed that human mental capacities are unique — they are of a different kind to other phenomenon. This line of thought has combined with the tendency to think that only adult, white males are properly *human*. Women, non-white people, children, people with cognitive disabilities and non-human animals have been mistreated as a consequence of such ideas.

Granting that mind begins in perception and that non-mammalian creatures, perhaps even certain arachnids, have perceptual capacities is a good vantage point for unravelling and understanding how human cognition is not radically different from other species' capacities. Perhaps mankind entertain species-specific mental abilities. There is, however, no evidence that these abilities must be of a different *kind* to the abilities of other species. My ambition in chapter 4 of this thesis has not been to suggest that arachnids, or other non-mammalian creatures that we do not certainly known to have experience, do not have intentional minds. If there is any evidence that they do, I believe we have *normative* reasons for assuming it. Rather, my suggestion would be that there are significant reasons to suppose that intentional mind is experiential. If this is the case, then, provided that the creatures I mention above have intentional perceptual capacities, they have experiences. An important moral is that perception, whatever it is, is a primitive ability, shared across species, and that this capacity serves as the foundation for much of our mentality in general. As mentioned, this mentality has been used as an excuse to treat those assumed to lack it, in ways inferior. The philosophy of mind *should* counteract such practice.

Literature

Block, N. 2007, 'Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience'. *Behavioural and Brain Sciences,* 30, 481–548

Burge, T. 2010, *Origins of Objectivity.* New York: Oxford University Press

——— 1979, 'Individualism and the Mental', Midwest Studies in Philosophy, vol. 4, iss 1, 73-121

Dennet, D. 1984, *The Intentional Stance.* Cambridge, Mass: MIT Press

Frege, G. 1982, 'On Sense and Nominatum'. In A.P. Martinich (ed.) *The Philosophy of Language.* 2010, Oxford: Oxford University Press, 5th ed, 217-279

Geisler, W. 2011, 'Contributions of Ideal Observer Theory to Vision Research'. *Vision Research 51, 771–781*

Geisler, W. & Diehl, R. L. 2003, 'A Bayesian Approach to the Evolution of perceptual and Cognitive Systems'. *Cognitive Science,* vol. 27, 379-402

Hatfield, G. 2002, 'Perception as Unconscious Inference', In Heyer, D. (ed.) *Perception and the Physical World: Psychological and Philosophical Problems in Perception.* 2002, New York: Wiley & Sons

Kriegel, U. 2013, 'The Phenomenal Intentionality Research Program,' in Kiregel, U. (ed) *Phenomenal Intentionality.* 2013, Oxford: Oxford University Press

——— 2005, 'Naturalizing Subjective Character'. *Philosophical and Phenomenological Research,* LXXI, July 2005

Ludwig, K. 1996, 'Why things Look the Way They Do', in Akins. K (ed.) *Perception.* 1996, New York: Oxford University Press

Nagel, T. 1974, 'What is it Like to be a Bat'. *The Philosophical Review* LXXXIII, 4, 435-50

——— 1986, *The View from Nowhere.* New York: Oxford University Press

Neander, K. 2012. 'Teleological Theories of Mental Content', Edward N. Zalta (ed.), Stanford Encyclopedia of Philosophy, spring 2012 edition. Availiable at: URL = <http://plato.stanford.edu/archives/spr2012/entries/content-teleological/>.

Philips, I. 2015, 'Consciousness and Criterion: On Block's Case for Unconscious Seeing'. *Philosophical and Phenomenological Research,* LLC, August 2015

Prinz, J. *When is Perception Conscious.* Available at: http://subcortex.com/

Putnam, H. 1975, 'The Meaning of "Meaning"', in K. Gunderson. (ed.), *Language, Mind and Knowledge:* Minnesota Studies in the Philosophy of Science VII. 1975, Minneapolis: Minnesota University Press

Rock, I. 1975, *An Introduction to Perception.* New York: Macmillan

Rock, I. 1975b, 'In Defense of Unconscious Perception', in Norman, D.A. & Rumelhart, D.E. (eds.), *Explorations in Cognition,* 1975, San Francisco: W.H. Freeman & Company

Rescorla, M. 2013, 'Bayesian Perceptual Psychology' in Matten, M (ed.), *Oxford Handbook of the Philosophy of Perception,* New York: Oxford University Press (print version forthcoming) available online at: http://www.philosophy.ucsb.edu/docs/faculty/papers/bayesian.pdf

Wade, N. & Swanton, M. 1991, *Visual Psychology: An Introduction.* London: Routledge