UNIVERSIDADE ESTADUAL DE CAMPINAS
INSTITUTO DE FILOSOFIA E CIÊNCIAS HUMANAS

Daniel Credico de Coimbra

REPRESENTATIONALISM IN PERCEPTUAL PSYCHOLOGY AS PER TYLER BURGE

REPRESENTACIONALISMO NA PSICOLOGIA PERCEPTUAL DE ACORDO COM TYLER BURGE

CAMPINAS
2022

Daniel Credico de Coimbra

REPRESENTATIONALISM IN PERCEPTUAL PSYCHOLOGY AS PER TYLER BURGE

REPRESENTACIONALISMO NA PSICOLOGIA PERCEPTUAL DE ACORDO COM TYLER BURGE

Dissertação apresentada ao Instituto de Filosofia e Ciências Humanas da Universidade Estadual de Campinas como parte dos requisitos exigidos para a obtenção do título de Mestre em FILOSOFIA.

Dissertation presented to the Institute of Philosophy and Human Sciences of the University of Campinas in partial fulfillment of the requirements for the degree of Master in PHILOSOPHY.

Orientador: Prof. Dr. Marco Antonio Caron Rufffino

ESTE TRABALHO CORRESPONDE À VERSÃO FINAL
DA DISSERTAÇÃO DEFENDIDA PELO ALUNO
DANIEL CREDICO DE COIMBRA, E ORIENTADA
PELO PROF. DR. MARCO ANTONIO CARON RUFFINO.

CAMPINAS

2022

Ficha catalográfica
Universidade Estadual de Campinas
Biblioteca do Instituto de Filosofia e Ciências Humanas
Cecília Maria Jorge Nicolau - CRB 8/3387

C665r

Coimbra, Daniel Credico de, 1996-
Representationalism in perceptual psychology as per Tyler Burge / Daniel Credico de Coimbra. – Campinas, SP : [s.n.], 2022.

Orientador: Marco Antonio Caron Ruffino.
Dissertação (mestrado) – Universidade Estadual de Campinas, Instituto de Filosofia e Ciências Humanas.

1. Representação (Filosofia). 2. Percepção (Filosofia). 3. Ciência - Filosofia. I. Ruffino, Marco, 1963-. II. Universidade Estadual de Campinas. Instituto de Filosofia e Ciências Humanas. III. Título.

Informações para Biblioteca Digital

Título em outro idioma: Representacionalismo na psicologia perceptual de acordo com Tyler Burge
Palavras-chave em inglês:
Representation (Philosophy)
Perception (Philosophy)
Science - Philosophy
Área de concentração: Filosofia
Titulação: Mestre em Filosofia
Banca examinadora:
Marco Antonio Caron Ruffino [Orientador]
Sérgio Farias de Souza Filho
Nara Miranda de Figueiredo
Data de defesa: 27-05-2022
Programa de Pós-Graduação: Filosofia

Identificação e informações acadêmicas do(a) aluno(a)
- ORCID do autor: 0000-0003-2084-748X
- Currículo Lattes do autor: http://lattes.cnpq.br/3480933478080935

UNIVERSIDADE ESTADUAL DE CAMPINAS

INSTITUTO DE FILOSOFIA E CIÊNCIAS HUMANAS

A Comissão Julgadora dos trabalhos de Defesa de Dissertação de Mestrado, composta pelos(as) Professores(as) Doutores(as) a seguir descritos, em sessão pública realizada em 27 de Maio de 2022, considerou o candidato Daniel Credico de Coimbra aprovado.

Prof. Dr. Marco Antonio Caron Ruffino

Dra. Nara Miranda de Figueiredo

Dr. Sérgio Farias de Souza Filho

A Ata de Defesa com as respectivas assinaturas dos membros encontra-se no SIGA/Sistema de Fluxo de Dissertações e na Coordenadoria do Programa de Pós-Graduação em Filosofia do Instituto de Filosofia e Ciências Humanas.

Acknowledgments

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

Agradecimentos

O presente trabalho foi realizado com apoio da Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Código de Financiamento 001.

Resumo
Em *Origins of Objectivity* (2010), Tyler Burge interpreta a psicologia da percepção. Ele articula claramente o que a ciência visa explicar e como suas explicações funcionam. Nós assumimos que a ciência possui as propriedades básicas que ele a alega ter; tais propriedades são tornadas explícitas. De acordo com ele, a ciência sistematicamente emprega em suas explicações uma noção de representação irredutivelmente normativa. A Seção I examina os termos, assunções, e formatos explicativos da versão de Burge da psicologia perceptual. A Seçaõ II revisa a literatura empírica sobre tal mecanismo, avaliando a adequação do sumário de Burge. O exato papel explicativo das constâncias é desenvolvido na Seção III, introduzindo a noção de explicação teleológica. A literatura filosófica sobre visões teleológias da representação é revisada na Seção IV, nos permitindo explicar por contraste a visão teleológica dissidente de Burge na Seção V. Nos avaliamos seus argumentos e finalizamos oferecendo uma contra-proposta.
Palavras-chave: Representação; explicação teleológica; filosofia da percepção.

Abstract

In *Origins of Objectivity* (2010), Tyler Burge interprets perceptual psychology. He states clearly what that science aims to explain and how its explanations go about. We assume that the science has the basic properties he claims it has; such properties are made explicit. According to him, the science systematically employs an irreducibly normative notion of representation in its explanations. Section I examines the basic terminology, assumptions, and explanatory formats in Burge's rendition of perceptual psychology, in which perceptual constancy takes center stage. Section II reviews the empirical literature on such mechanism, assessing the accuracy of Burge's summary. The exact explanatory role of constancies is worked out in Section III, introducing the notion of a teleological explanation. The philosophical literature on teleological views on representation is reviewed in Section IV, paving the way for us to contrast them with Burge's dissenting teleological view in Section V. We evaluate his arguments and round up by offering a counter-proposal.

Keywords: Representation; teleological explanation; philosophy of perception.

TABLE OF CONTENTS

Preface

What the Philosophy of Perception is

Perception is a major topic in epistemology and metaphysics. It presents a continuous source of puzzlement for researchers and has elicited multiple different approaches in its resolution. The present dissertation is an exploration in the philosophy of perception, and thus has to grapple with this difficult subject.

We suppose that it is proper to begin by illustrating three main reasons why perception continues to puzzle humanity. The first reason is that perception bundles together thorny semantic, epistemic, and ontic questions. What we mean is that perception can be studied as a way to acquire reference and concepts (semantics), as a way to acquire information and knowledge (epistemology), and as an entity or process itself (ontology). These three ways of studying perception may not be independent; answers to one inquiry may provide insight into another.

The second reason is that the parties discussing perception use the same terms—‘reference’, ‘concept’, ‘content’, ‘representation’, ‘semantics’, ‘information’, ‘knowledge’, ‘access’, ‘directness’, ‘epistemology’, ‘entity’, ‘state,’ ‘process’, ‘experience’, ‘sensation’, ‘ontology’, —with different and sometimes undisclosed meanings. To avoid a theory-laden exposition, below we use these terms generically. We will render them more precise as it becomes relevant in the dissertation.

The third reason is that perception is multifarious. We perceive our bodies, our environment, and perhaps our mental processes. In some sense we perceive a skin cut, room temperature, and perhaps feelings such as hunger and surprise. Possibly in some other sense, we also perceive whirlwinds moving leaves, other people’s mental states moving their bodies,

as well as causal relations transmitting momentum between bodies. Researchers have long asked whether we perceive all these things in the same way: is there not a fundamental type distinction between how one perceives one's phenomenology, one's body, one's immediate surroundings, and "invisibles" such as whirlwinds, others people's mental states, and causal relations?

There might be, and the distinctions between those types have consequences that may concomitantly cut across epistemology, semantics, and metaphysics. This is strongly highlighted in the discussion about direct perception. Some contend that we can directly perceive only our phenomenological events. Such folks *tend* to understand perceptual states as phenomenological events themselves (metaphysics) and furthermore as non-representational in content (semantics); and those two have natural connections to the view that we have merely inferential and possibly dubious access to the external world (epistemology).

Sweeping over many subtle intermediate positions, we reach at the other extreme replying that we can in fact directly perceive things as external and "pictorially invisible" as causal relations in external mechanical processes. This position is *often* accompanied by the view that perceptual states are those which have obtained information about the world in a certain way (metaphysics, semantics), which in turn has natural connections to epistemic reliabilism and the denial that access to the external world requires inference, for one or another sense of 'inference' (epistemology).

As the two above examples indicate, the three main philosophical areas intertwine. Keep in mind that each such view also affects wider arenas in philosophy; for instance, consider how much hinges (or has been thought to hinge) on the correct theory of concept acquisition, and how much *that* seemingly hinges on what kinds of content perception turns out to have. The philosophy of perception threads over multiple fundamental issues in way that ends up tying it to positions on things other than perception itself.

Should views in semantics, epistemology, and metaphysics necessarily go together when discussing perception, as they artifactually did in the two archetypes constructed above, philosophy of perception would have a tidier landscape. Unfortunately, the fact is that with sufficient cleverness one can tease them apart and make them vary somewhat independently. For example, senior researchers in skepticism might view their topic as independent from questions about direct *versus* indirect perception or internalism *versus* externalism on the semantics of perception. So there is even contention about in what ways these different areas connect.

What's more, different views in each philosophical area lead researchers to define the key terms mentioned above in different ways; only painstaking attention to detail allows one to use the terms in theory-neutral ways. For instance, some among those who deny perception is *representational* still state that it has *content*; as a result, we should not construe questions about content as necessarily *semantical*. But this is only insofar as semantics is that which has to do with representationality, which is, in fact, another point of radical contention. And so on, to the point that occasionally one finds it expedient to swirl at length around fine points of terminology and classification, which is useful enough, but which does prevent attending to more substantial brawls.

To sum up, perception is puzzling chiefly because it bears upon and bundles together fundamental issues from multiple areas, making it internally complex and externally constrained by results outside the philosophy of perception. Progress, although very significant, has been hampered because philosophers have (understandably) failed to find a sufficiently abstract and neutral language in which to formulate problems and solutions.

During our dissertation, we have aimed to make our assumptions and terminology clear.

Introduction

The Dissertation Plan

The present dissertation is a commentary on a book on the philosophy of perception, *Origins of Objectivity* (Tyler Burge, 2010). We intend to articulate and critically evaluate that researcher's position on perception. Our criticism will be formulated without straddling too far from his basic presuppositions. This is because navigating the topic is made significantly easier, and in fact more profitable in a certain way, by holding fixed certain assumptions and arbitrarily choosing a standard terminology. For the sake of tractability, our work below will explicitly articulate and also *accept* many terminological and theoretical decisions laid out in Burge's work, and it is based on such terminology and assumptions that we will evaluate his main proposals. As such, our criticism of Burge's system of thought is offered in a *reformist* rather than a *revolutionary* spirit, if you will.

It is crucial to understand that this book is a work in naturalistic philosophy, insofar as it appeals to results in a science known as *perceptual psychology* as partial justification for philosophical positions. That is, many of the decisions behind his programme are based on his interpretation of perceptual psychology. Furthermore, the entire book upholds the presumption that empirical results can override apriori intuitions. For instance, the science is claimed to entail that perceptual representations exist, and his main defense of representationalism in the philosophy of perception is bolstered in this understanding of perceptual psychology. Apriori qualms against representations would therefore have to be set aside.

His work is part of a wider trend in philosophy of perception, which has seen a gradual increase in empirically-informed work. This is not, however, to say that philosophers have gotten closer to a consensus on key questions about perception. Empirical research on perception is split into diverse and possibly incompatible paradigms. Perceptual psychology is

located within the mainstream research paradigm, known as ‘cognitive science’, but the paradigm itself is hotly contested. This means that other philosophers might disagree with Burge on the grounds that perceptual psychology is incompatible with other scientific avenues whose research seems more plausible to them. Examples include what are called *enactive* (“action-based”) accounts of perception, which claim *not* to employ notions of representation. Whereas some varieties of action-based accounts intend to account for but a *segment* of perceptual phenomena, in effect *supplementing* work in perceptual psychology, other enactive theories aim to be all-encompassing and would thereby seem incompatible with Burge’s preferred models.

What he claims is that perceptual psychology is much more developed than other philosophers and scientific dissidents usually acknowledge. For one, it has achieved significant integration with other emerging empirical sciences: evolutionary and cognitive ethology (282, 319ff),[1] human developmental psychology (xiv, 284),[2] physiological and psycho-physical optics (272, 383) and, under some approaches, with computational neuroscience (see Frisby & Stone, 2010).[3] Furthermore, he claims, perceptual psychology has developed detailed empirical models for many aspects of animal, infant, and adult human perception (87, 342-66). Their surprising predictions have been often confirmed and their assumptions have led to fruitful research questions.

Finally, the science is claimed to be *explanatorily unified*, insofar as these models share their explanatory scheme—namely, the notion that perceptual states are external-world representations. Figuring out what such representations are, and what their explanatory role is, comprise the major interpretation efforts of *Origins of Objectivity* relative to perceptual psychology.

Herein comes our reformist approach to *Origins*. We will take for granted that perceptual psychology is the most well-developed scientific approach to perception. We will also, for the most part, accept Burge’s interpretation of perceptual psychology. Questioning

---

1 Respectively, the evolutionary and mentalistic studies of natural animal behavior; the connection is unsurprising given how perception guides behavior.

2 The mentalistic study of how adult emotive and cognitive capacities are developed from earlier cognitive capacities; this is relevant because as perceptual abilities are both developed through life as well as relevant to acquiring other abilities.

3 These latter three comprise the study of the mechanistic and computational properties of sensations and perceptual processing, and how these properties relate to the resulting perceptual states.

these two points would require obtaining systematic familiarity with the general field, an undertaking well beyond the scope of this dissertation.

As a result, we will take for granted his broad understanding of what perceptual psychology *aims* to explain and *how* it goes about explaining it. Our overall aim is to articulate the reasoning that led to the book's conclusion on those two major subjects.

Our work offers two major critical efforts. The first is directed at Burge's interpretation of *perceptual constancy*, a neuropsychological mechanism on which his conclusions heavily rely. The alleged mechanism of perceptual constancy assumes such a center stage that we will delve into the empirical literature behind it and question whether his interpretation is justifiable.

The second is directed at the philosophical consequences he derives from his interpretation of the science. Chiefly, we will argue that a family of positions known collectively as *teleofunctionalism* are compatible with perceptual psychology, contrary to his conclusion.

Like most other academic works, we do not intend to start from a fresh beginning. We examine what can be described as a conditional question: *if* perceptual psychology were true, specifically Burge's version of it, what consequences would flow to the philosophy of perception?

Part One

Perceptual Constancy

Origins of Objectivity has found it instructive to frame perceptual psychology fundamentally in terms of its explanatory targets. What a science wishes to explain *constrains* what explanatory forms are available to it and what theoretical posits it is forced to make. Based on Burge (87-98), this section outlines such explanatory targets and the explanatory form perceptual psychology developed to meet them. This will lead us to direct close scrutiny on the center-stage notion of *perceptual constancy*.

Section I – *Tyler Burge on Explanation in Perceptual Psychology*

Perceptual psychology centrally investigates perceptual states' *formation processes*. We will explain what a formation process is, how percentual constancies play a role in it, and the way in which perceptual psychology constructs explanations and type-categorizations based on perceptual constancies. Our exposition is based on Burge's outlining of perceptual psychology explanatory structures (cf. Burge 2010: 87ff).

A formation process is an internal process that begins when any bodily sensor is activated by an external stimulus. This activation constitutes what is called a *surface sensory stimulus*. The causal chain begins outside the body, before the formation process initiates. It begins with an external object, called the *distal stimulus*, initiating a causal chain (e.g. emitting photons, radiating heat, generating acoustic waves) that ends up altering the state of an organism's detectors (e.g. the retina, the skin, the eardrum). Such alteration is what we have called a surface sensory stimulus. It is also referred to as *proximal stimulus*. These processes leading up to proximal stimuli are the province of optics and other physical sciences (Burge 2010: 89).

Perceptual psychology is concerned with what comes afterwards, namely, the effects that surface, proximal stimuli exert in the organism's neurocircuitry so as to produce perceptual states (Burge 2010: 89-90). For instance, the inner mechanisms that allow animals to visually or olfactorily recognize their kins, allies, and foes within a community.

There are fascinating problems to be solved in the domain of perceptual processing. The most general is to describe the process such that the description employs categories that allow for law-like generalizations, themselves hallmarks of systematic scientific explanation (Burge, 2010: 49-51). We say 'law-like' because there are no *exceptionless* generalizations about an organism's messy functioning.

In other words, the general aim is to carve up brain and body functioning in chunks which are regularly and systematically involved in producing various forms of perceptual states. Such chunks may be low-level, as exemplified when vision neurophysiologists study the retina's computational procedures, as well as high-level, as happens when researchers invoke concepts from systems neuroscience and representational theories of mind.

According to Burge, perceptual psychology has described animal and human perceptual formation process by highlighting a mechanism with great philosophical consequences: *perceptual constancy*, a mechanism he exemplifies extensively (cf. Burge, 2010: 351ff; 408ff). Perceptual constancy is the process of intaking very *distinct* proximal stimuli—such as retinal patterns and eardrum vibrations—and outputting very *similar* perceptual states. As such, perceptual constancies, quite literally, *group* different proximal stimuli by treating them alike (Burge, 2010: 114; 274; 408).

The need for perceptual constancy stems from the messiness of proximal stimuli. Realistic painters and machine vision researchers are painfully aware that even an object as simple as a featureless monochrome wall can reflect a hundred subtly distinct shades of color even under uniform white illumination, a problem made an order of magnitude worse when surface irregularities are anywhere introduced. Similar issues torment audio producers and machine hearing researchers: the very same phoneme can have wildly distinct acoustic profiles under different contexts, not to mention all the background acoustic noise present in any real-world situation. An everyday example, now involving high-level perception, is how we can perceive persons as having the same faces over time despite seeing them from different angles and under different illuminations. This phenomenon, however, shows itself even for things as simple as shape and color detection in early visual processes (see our Section II for empirical examples).

Perceptual psychologists, Burge claims, have found that perceptual constancies are involved in most law-like generalizations required to understand the perceptual system in animals and humans. It would, however, befuddle researchers in case the identified perceptual constancies performed grouping operations that, in some sense, "made no sense". Constancy operations make sense when they effectively ensure that *similar environmental causes* induce *similar perceptual states*. (Burge, 2010: 87ff; 233-4; 260; 274-6).

For example, it could group all acoustic patterns regularly associated with a linguistic phoneme—no matter how distinct they might be—without including patterns *not* associated with it. That is, it would make sense to produce the same perceptual state in response to the same phoneme. The result is a perceptual constancy that makes sense: the same *distal* phoneme, although inducing distinct *proximal* stimuli in an organism, leads it to have the same perceptual state.

As it turns out, that is exactly the way in which perceptual psychologists have made sense of perceptual constancy, according to Burge (see references above). The proximal stimuli in each grouping have no unifying characteristic—certainly not in their physical make-up, given how distinct they are—*except* that they are regularly caused by the same distal cause. As such, each perceptual constancy mechanism groups a set of proximal stimuli just in case they are regularly caused by the same distal cause. We can refer to this as a *matching* or a *mirroring* process.

Perceptual psychologists, he goes on to claim, were therefore led to categorize perceptual constancy mechanisms based on what distal cause they mirror—that is, what distal cause unifies the proximal stimuli which they happen to group. We might as well use 'perceptual constancy' to refer only to grouping mechanisms that match some distal cause.

What is more, on his reading, perceptual psychologists categorize perceptual states themselves based on what perceptual constancy mechanism was involved in producing them. It is this move which gives perceptual constancy momentous philosophical consequence. We will end this section showing how perceptual constancy figures in the categorization of perceptual states and, what is more, in the characterization of perception as representational, thus motivating the relevance of the empirical review we perform in the follow-up section.

Any scientific theory carves up reality into distinct kinds or, equivalently here, types. The way in which reality is carved up might be called the science's *typing system*. Typing systems are useful to determine which are the relevant details in a domain. In a typing system, there will be specified criteria for two objects to be grouped together and then treated as

equivalent or at least very similar, despite their differences. These groups are what are called *types*. Typing systems allow for wider generalizations. They are wider because they refer to more abstract patterns in a domain; the patterns are more abstract due to ignoring some individual differences. Every scientific field employs its own typing system to systematize its understanding of the world.

Illustrative examples are found within the legal system. Law's central task is not truth, but to create a workable system with useful incentives, fair compensation, and fair punishment. That task is pragmatically tractable only by intelligently grouping ("codifying") messy human activities into more neatly distinguished types. For instance, *tax evasion* is a broad type which encompasses quite different activities—ranging from smuggling and money laundering to hiding assets and underreporting tax liability—but which are treated as quite similar when it comes to their legal consequences. More specific categories, such as *money laundering*, themselves can take a multitude of forms, but their instances are treated even more similarly. There is, however, some degree of unnaturalness to Law's basic categories, since the objects grouped together may not in fact be similar enough.

Still, the above examples bear a structural analogy to a truth-seeking science. This is because the latter aims to organize a domain which is also messy in its details. The difference is that systematic science is possible because that domain is, as a matter of fact, quite regular and systematic under a certain level of abstraction. At least from a scientific realist perspective, there are more or less objective ways to carve up reality, and that is what science aims to do and, to various degrees, actually achieves.

A typing system is *accurate* when its groupings match the domain's objective kinds. To speak metaphorically, it must carve the world in the same way it is naturally carved. For instance, biologists seek to identify the natural units that are most relevantly involved in interesting biological patterns. This has led biology to create unifying abstract categories such as *blood*, *cell*, *organ*, *skin*, *tropism*, *wing*, *adaptation*, *cell differentiation*, *homeostasis*, *metabolism*, *reproduction*, and so forth; categories which group quite different instances coming from quite different animals, because such instances share very relevant similarities.

Origins of Objectivity is committed to scientific realism and so are we, in reformist spirit.

One of its tenets is that sciences are committed to accuracy, as mentioned above. A scientific theory's typing systems, therefore, are intended to be accurate.

When do we have good evidence for scientific accuracy? The answer comes from another tenet of scientific realism: that a science's *sucess* is our most reliable guide to its accuracy. Successful sciences are those with reliable predictions, stable explanatory models, unifying capacities, and that lead to interesting new questions. This kind of science is also regularly referred to as *fruitful*.

It is Burge's understanding that perceptual psychology has achieved a strong degree of success. As we have mentioned in the introduction, that science developed precise models for multiple aspects of perception, each of which has led to new insights and surprising predictions, and all of which have employed the same underlying explanatory scheme.

For example, the many models required to achieve a panoramic understanding of vision *all* make the assumption that the visual system is ingrained with the *assumptions* about the external world—and the very same assumptions, to boot—and that the visual system employs perceptual constancy mechanisms to discriminate distal enviromental conditions. In more detail:

> *"The dominant scheme in the psychology of vision for explaining how [underdetermination] problems are overcome goes back to Helmholtz. The idea is to explain a series of unconscious, largely automatic transformational processes that lead from registration of the array and spectral properties of light striking the retina to the formation of perceptions as of specific aspects of the distal environment. The transformations operate under certain principles that describe psychological laws of law-like patterns. These laws or law-like processes serve to privilege certain among the possible environmental causes over others. The net effect of the privileging is to make the underdetermining proximal stimulation trigger a perceptual state that represents the distal cause to be, in most cases, exactly one of the many possible distal causes that are compatible with (but not determined by) the given proximal stimulation. I call psychological principles that describe, in an explantory way, these laws of law-like patterns* formation principles*."*
>
> — Burge (2010: 92)

These models about vision explain the conditions under which we perceive veridically and unveridically in a quite systematic way. They have also made predictions about under what precise conditions humans should undergo illusions, predictions which have been confirmed

in further experiment. Should all this be so, there would be good evidence that the typing systems which figure in perceptual psychology are, in fact, accurate.

Perceptual psychology characterizes what it is to be a *perceptual state* through a few typing systems. That is, perceptual states acquire their basic individuation criteria based on types within these systems. Should two entities be grouped together, they will be considered to be the same perceptual state. Should an entity not fall under any categorization, it will not be considered a perceptual state. To be sorted into different or identical types in a perceptual typing system is, respectively, to be different or type-identical perceptual representations.

It is Burge's understanding that these typing systems figure *ineliminably* in the science. That is, they are not mere ways of speaking that facilitate teaching and research, but are rather central parts to the science's predictive and explanatory models.

These typing systems have three fascinating and interconnected aspects we will explore, according to Bunge. The first is that perceptual kinds are individuated by perceptual constancy mechanisms, the topic of this dissertation's first half. The second is that, due to the characteristics of a perceptual constancy mechanism, perceptual kinds constitute *representational kinds*. The third is that perceptual kinds are defined through *teleology*, a claim also related to perceptual kinds being representational kinds. These latter two aspects comprise the topic of this dissertation's second half. That a well-established science vindicates representations and teleologies—that is the momentous Burgean conclusion we explain and evaluate throughout this dissertation.

When it comes to perceptual states, Burge says, there is a rigid hierarchical typing system. There is a set of basic perceptual kinds we can call 'Ψ'. Its perceptual kinds are individuated based on a single basic characteristic, namely, the *type attribute* that is being perceived (Burge, 2010: 99; 296; 308-9).

Attributes, in Burge's terminology, comprise *properties*, *relations*, and *natural kinds* (Burge, 2010: 24n11). Perceptual states are thus characterized as being perceptions as of things such as a shape or color (property), parthood or a distance (relation), and water or physical body (kind). The perception as of type attributes characterizes basic perceptual representations.

Perceptual states can be further classified into sub-types. One example is the perception as of *token* attributes as well as *particular* objects. Another example is the perception as of an attribute *from a perspective*, as for instance a slanted disc. Perceptual psychology systematically investigates these sub-types, but they do not figure in Ψ and are

not relevant for the specific Burgean claims we will explore. A scientific-cum-philosophical study of such aspects of perception can be found in Burge (2005).

The basic typing system used in perceptual psychology, which we have called Ψ, organizes *internal sensory states* into basic *perceptual states*. That is, of all the things in the world, only internal sensory states can fit the conditions to belong to a type in Ψ. Plus, not all sensory states fit such conditions, and therefore the science distinguishes between merely sensory states and properly perceptual states (Burge, 2010: 9; 263-5; 316). Internal sensory states are effectively *re-grouped* Ψ: different sensory state types will be grouped together, and tokens of the same sensory state type might be grouped separately.[4]

Sensory states comprise both the surface registrations we have called 'proximal stimuli', such as retinal activation patterns, and internal neurological states that result from processing proximal stimuli. It is not relevant to our purposes to go much further into what sensory states are; the relevant details can be found in Burge (2010: 315-9, 376-8, 421-30).

As stated, only internal states can satisfy the conditions to constitute perceptual states. The reason is that those conditions are formulated in terms of an *internal formation process*, already mentioned above, and of which surface stimuli have none. As seen earlier, such formation process must include a perceptual constancy mechanism, a process that bears much weight under Burge's analysis. The constitutive role of inner processing in individuating perceptual states is expressed in the following passage:

> *"Nature molds all sensory systems perceptual and non-perceptual to be likely to respond to environmental conditions that are beneficial to animals' functions. Perceptual systems are distinctive in the way that this likelihood is determined [namely, through perceptual constancy]. The beginning of understanding this way lies in reflecting on the formation principles. (...) The formation principles describe laws of perception formation that reflect constitutive determination of perceptual representational kinds by the distal regularities and kinds. They illustrate anti-individualism regarding perception."* — Burge (2010: 345-6)

4 We use the term 'state' as a generic term for states, events, processes, or instances of any other ontological category which may be involved in our sensorial interaction to the world. We are not at the stage of theoretical development where much precision can be afforded to such terms, much like someone discussing scientific theories needn't decide on whether such theories are a set of sentences, of propositions, of counterfactuals, or some other thing.

To sum up, perceptual kinds are fundamentally characterized as a representation as of a certain type attribute, and the type attribute associated to a perceptual state is constitutively connected to the type attribute mirrored by the perceptual constancy mechanism that produced the perceptual state.

More rigorously, an internal sensory state constitutes a perceptual representation as of type attribute *A* if and only if it is caused, with law-like regularity, by a perceptual constancy mechanism triggered by the distal cause *A* (Burge, 2010: 345-7). For instance, suppose that sensory state *S* virtually always owes its existence to *redness* tokens, whose proximal impact on the organism is processed by a color-discriminating perceptual constancy mechanism. Then *S* is a perceptual representation as of redness.

That is not all. In addition, the perceptual constancy mechanism must have a *teleology* to detect *A* tokens (Burge, 2010: 75-6; 81; 309-10). Exploring what a teleology is, and what motivates Burge's claim that representations are related to teleology, will be fully carried out in the second part of this dissertation.

For now, we will investigate Burge's claims on perceptual constancy. We ask the following questions: Is perceptual constancy as ubiquitous as he claims? Do they take center-stage in the definition of perceptual kinds? And, finally, are they described using the normative, teleological, goal-oriented language of purposes?

Section II – *The Science of Perceptual Constancy – Empirical Review*

We have highlighted the centrality taken by perceptual constancy in Tyler Burge's recent work on perception, with ramifications for semantics (e.g. Burge, 2009b), epistemology (e.g. Burge, 2003a, 2003b, 2005: §VII), the philosophy of action (e.g. Burge, 2009a), and the nature of mentality. He takes it to be the source of objective perceptual representation, that is, perceptual reference to external, mind-independent particulars, properties, and relations. In this section, we will cover in more detail how scientists have come to understand perceptual constancy and see it applied in a few detailed cases, a process during which we examine competing frameworks for understanding the issue. We finish by listing some concerns that have been raised in the empirical and philosophical literature about the very concept of perceptual constancy and its philosophical usefulness.

Theorization about perceptual constancy, either in general or applied to a specific domain, has not yet reached scientific convergence. There are diverging frameworks for

understanding how peceptual constancy works fundamentally, as shown by the theory taxonomy in Shebilske & Peters (1996: §§2-4). That is, there is still no unified framework to understand how the brain filters the perspectival and idiosyncratic features of proximal stimuli to arrive at the invariant, observer-independent features of external reality.

A larger proportion of theories understand that perceptual constancy is the product of complex unconscious processing, going as far back as Hermann von Helmholtz in the 19th century. Some have argued that such approaches are ultimately untestable due to the indefinitely malleable nature of the posited hidden mechanisms (cf. Wilcox 1992: §2 and more recently Wagner 2012: 74-6). Still, it has produced a very large field with precise and predictive explanations for uncountably many aspects of visual perception, with long textbooks devoted to introducing the field to the gifted amateur (Frisby & Stone 2010).

Some researchers have offered what is called a *relational* approach to constancy, either as a result of a wholesale rejection of unconscious processing models, or merely by believing that such an approach is more adequate for a circumscribed set of cases. This approach undersands that the type attributes related to perceptual constancies are correlated with higher-order invariants in proximal stimuli, such as an object's relative size or brightness to its surroundings as an observer moves away from it or incident illumination changes (Shebilske & Peters 1996: 229-31). No internal processing is required to mirror external type attributes because they are already directly tracked by patterns in surface sensory stimulus.

However, Burge (2012: 91-4) describes perceptual constancy as requiring complex internal processing. The reason is that achieving sensory constancy in any other way would involve a mechanism so simple that it could be described without teleological terms, which, as we shall see in this dissertation's latter half, are what mark the distinction between what is merely sensory and what is truly perceptual. On Burge's construal, then, relational theories of perceptual constancy are either erroneous or describing what he would call a merely *sensory* constancy.

Another point of theoretical dissent was on how to categorize the radical stimulus variation across which we exhibit perceptual constancy. Many early theorists regarded it to be variation in *conscious experience* (Schulte 2021: §2, Shebilske & Peters 1996: §3.1). This marks a fundamental departure from how Burge and most recent theorists view perceptual constancy (Schulte 2021: §2). To them, proximal stimuli are surface sensory stimuli, and not

conscious experiences. This is so specially because perceptual constancy does not require conscious experience, as demonstrated by the fact that blindsight and a few other conditions fulfill all operational criteria for perceptual constancy without being associated to conscious experience (Burge 2010: 374-6).

Even within a particular framework for understanding perceptual constancy, such as one considering internal processing based on surface sensory stimuli, there is large-scale disagreement on how the constancy mechanism is implemented. This is well exemplified by the existence of two prominent, and opposing, classes of models for color constancy listed in Brown (2003: 256-8), both of which Brown himself believes to be fundamentally inadequate.

Vision is the most studied perceptual faculty (Burge 2010: 88-90, 420), and color constancy is among the most studied phenomena of perception, and there has yet to be scientific convergence relative to its nature (Brown 2003: 268-9; Hoffman 2003a: 374; Hoffman 2003b: 435), much as convergence has not been reached after over a century of research into size constancy (Wagner 2012: 75).

Burge's commentary on perceptual constancy, however, glosses over any such scientific dissent. His interpretation would seem to be broad enough to apparently avoid commitment to any specific model or account, except the aforementioned commitment to internal processing based on surface sensory stimuli. This suggests that the philosophical consequences he derives are intended as account-neutral, unless a significant revolution occurs in the field.

This might be so, for there apparently *are* broad, account-neutral remarks that can be made about perceptual constancy, and which might be enough to justify his conclusions. After the above overture reviewing dissent relative to perceptual constancy, let us proceed to rendering explicit what is seemingly common ground for perceptual constancy researchers.

As we have already noted, nowadays it is a near consensus that perceptual constancy is the ability to "recognize", in some disputed sense, some shared *environmental* cause amongst *radically distinct* proximal stimuli. This is a notable example of scientific convergence relative to perceptual constancy (Schulte 2021: §2).

The above account-neutral remark is already by itself helpful to deriving philosophical consequences. We will give one example. *If* one accepts that perceptual constancies for a type attribute can establish perceptual reference to that type attribute, *then* the aforementioned

remark is sufficient to solve what has been called the *distality problem*. That problem is to answer the following challenge: given that multiple entities are involved in the process that goes from an external environmental source all the way to the resulting internal percept, what would determine whether our perceptual reference is established to one or another among the elements in that causal chain? (Roche & Sober 2019: §§2.2, Schulte 2021: §3).

The problem would be solved, we argue, because we are not said in perceptual psychology to bear perceptual constancy relative to photons, retinal patterns, the optical nerve, or other components in the causal chain linking the distal cause to the resulting percept. And the reason for that is that those intermediary components can vary radically while the connection between the distal cause and the resulting percept remains the same. The perceptual mechanism in question is characterized by constancy across variations in these intermediary steps. It is for these reason that, if perceptual constancy establishes reference, then we refer to external entities at the end of the causal chain rather than any intermediary steps. This conclusion has been drawn by Tyler Burge (e.g. Burge 2010) and Peter Schulte (e.g. Schulte 2018: §5-7).

Animals and humans exhibit perceptual constancy with respect to numerous environmental properties, including color, distance, location, mechanics (e.g., rigidity and texture), movement direction, object edges, object identity, orientation, relative position (e.g., foreground/background relations), shape, size, speed (linear and angular), and temperature. Under the account in which perceptual constancy establishes perceptual reference, such external properties and attributes enter into animal and human mental content. None of these include retinal or other intermediary properties. Perceptual constancy is distality-oriented.

Scientific convergence has also been reached in the recognition that discriminating distal feature is a *remarkable* feat whose attainment requires the perceptual system to be *attuned* to environmental regularities. To illustrate this fact, we will delve further into how perceptual constancy works and examine a few examples more closely. It is important to understand how perceptual constancy works in order to grasp what is necessary for it to possibly work.

Perceptual constancy is remarkable because proximal stimuli strongly *underdetermine* their distal environmental causes, in the sense that multiple environmental sources could have caused the same surface sensory effects. Worded differently, every type proximal stimulus—

every pattern of surface sensory activation—can in principle be caused by radically distinct type atributes. This is true for distal features as simple as a surface's reflectance properties (MacLeod & Golz 2003: 209), and it in fact generalizes to multiple if not all aspects of vision and in fact perception in general (Burge 2010: 87-92).

The underdetermination problem can be framed into two components. One component is that identical objects can seem different under different conditions. On the one hand, and thankfully, objects bear properties that are invariant across space, time, illumination, as well as perspectival variations such as angle and distance. On the other hand, and quite unfortunately, our proximal stimuli are radically *variant* across these changes. The very same object strikes us very differently under different conditions. As such, environmental invariances are *masked* by contextual and perspectival idiosyncrasies, which are reflected in proximal stimuli. Were it not for perceptual constancy—i.e., if differences in proximal stimuli were taken at face value—, identical objects and properties would not be recognized as the same, making perception faulty.

The other component is that distinct objects can produce identical proximal stimuli, for perception involves loss of information. A major example is that vision involves projecting a three-dimensional environment into a two-dimensional retina, which leads to a mathematically necessary loss of information (Pinker 1997: 9-10, 215-8). Multiple three-dimensional patterns could have lead to the same two-dimensional retinal pattern. Similar remarks hold for aspects of visual perception in general (Burge 2010: 90, Pinker 1997: 28-9). In fact, as we have stated, scientists working in perception have repeatedly pointed out how every perceptual modality faces this predicament in some form.

Here are a few examples. Under the same illumination, paper sheets and coal lumps imprint the retina very differently. The former is pitch black and the latter is bleach white. Yet, coal lumps can be thousands of times *brighter* than paper sheets, should the former be placed under the sun and the latter be placed in a living room. The perceptual system is attuned to the fact that how bright an object looks (its "luminance") is a product of two factors: how bright a light is shined upon it (its "illuminance") and how the object itself reflects light (its "reflectance", including color and shininess). By attending to the whole illuminance context, our visual system can somehow overcome the Sun-induced brightness mask and figure out that the coal lump is in fact pitch black.

A second example is color constancy. A constancy for redness is the ability to detect that an object is red across multiple different illumination conditions, including different wavelengths and cone activation patterns. It also involves discerning when an object is *not* red even when it appears to be so. For instance, shining monochromatic red light over a plain blue object will make the object reflect a lot of red light and thus appear to be red; perceptual constancy fails in this artificial scenario. Yet, on most occasions, multiple clues will allow our visual system to figure out that the illumination conditions are misleading. For instance, the whole scene will be uncommonly red, leading the visual system to correct for the excessive redness. Our visual system takes significant advantage of cues provided from an object's environment, and it is more competent at doing this the more the environment approximates a well-lit natural scene, as exemplified by our increased ability for size constancy under such scenarios (Wagner, 2012: 63, 66).

A third example pertains to shape or size detection. Suppose we write a dumb shape-detection or size-detection piece of software in some programming language. It determines an object's shape or size merely by computing the shape or size of the image it produces in a camera's digital sensors.

If we feed the software a video of a rotating coin, the shape program will conclude that the coin is transforming from round to highly elliptical to round back again. Animals and humans are not so dumb: they realize that the coin has not changed shape at all. We are generally not fooled by changes that occur only in the retina. Plates on a dining table seem round even if we do not see them from a bird's-eye view; receding objects do not seem to shrink in tandem with their retinal projection; and so forth.

Similar considerations would hold for the size detection program. Receding objects would seem to shrink simply because their projection into the camera's sensors are shrinking, due to simple optical laws. We know better: receding objects seem to be constant in size, and we are not surprised when we reach toward far away (and retinally miniscule) objects and realize that all the while they have been the same size as our (retinally gigantic) hands.

Reaching those seemingly obvious statements requires a smartly-designed perceptual system. In the rotating coin example, the system must be attuned to the fact that *slanted* round disks are much more common than truly ellipsoidal objects, so a retinal ellipse probably matches onto an external slanted disk. Plus, objects usually do not change shape, which makes it easier to know that the coin is *rotating* rather than *shapeshifting*. In the receding object example, our size-detection apparatus is in tune with the spatial invariance of size as

well as the optics of perspective, and thus corrects for shrinking or enlarging retinal images. Our ability to judge absolute and relative size at a static distance is remarkable as well, although our performance drops considerably when objects are observed isolated from an environmental which could furnish perspectival cues; we mostly end up underestimating or overestimating relative size of objects at different distances from us (Wagner 2012: Table 3.1).

To sum up, differences in proximal stimuli do not entail differences in their environmental causes, and sameness in proximal stimuli does not entail sameness in their environmental causes. The fact that proximal stimuli *underdetermines* distal simuli is problematic because perception is strongly guided by proximal stimuli.

Fortunately, the environment contains law-like regularities to which our perceptual system are attuned, and which thus help us to overcome sensory underdetermination and to succeed in discerning ("detecting") when a specific external affair is present. That is, a stimulus-based process systematically detecting external affairs is only possible thanks to certain law-like regularities in the organism's environment. Perception is attuned to environmental regularities in such a way that allows it to systematically discriminate in the proximal stimuli what, on the one hand, is due to perspectival and contextual masking and what, on the other hand, bears genuine information about external invariant properties.

An example of an environmental regularity can be illuminating. Human faces cause certain retinal stimuli, depending on illumination and perspective conditions. The very same stimuli can be caused by confounding factors, such as photo-realistic paintings, virtual reality headsets, carefully arranged debris, and serendipitous shadows or rock formations. It is a running gag that many people sight Christ's face on toasts, clouds, and Martian hills. Perceptual recognition of human faces is made radically easier by the fact that, in the natural environment, it is exceedingly rare for anything except a human face to cause such retinal stimuli. The confounding contraptions imagined above are ecologically abnormal. Amongst common natural objects, human faces bear genuinely *unique* optical properties; they imprint our retina like nothing else does.

Burge expresses the above fact by saying that human faces constitute the *normal environmental cause* of those retinal activation patterns, while other distal stimuli, such as rock formations, would constitute *abnormal* environmental causes (Burge, 2010: 98, 377).

The notion of a 'normal' cause makes biological sense due to law-like environmental regularities. As an example, the natural environment has a regular light source: a white Sun

shining through the atmosphere, generating illumination conditions that include mostly yellow, orange, and blue all reaching from above.

In the case of vision, objects rarely share their optical properties: their color distribution, their surface orientations, their shape and size, the way they reflect and diffuse light. This helps perserve the uniqueness of normal environmental causes, to which we must be attuned. A varied exploration of this fact can be found in Pinker (1997: ch. 4).

Plus, the fact that objects move, and that organisms too can move around to see stationary objects from different angles and distances in the so-called *perception-action cycle*, also helps eliminate a variety of potential confounders (Briscoe & Grush 2015: Introduction), further narrowing down the scope of normal environmental causes.

In summary, in normal environments, there can be sameness in distal stimulus (environmental cause) despite radical variation in proximal stimulus, due to perspectival and contextual variations, *but* sameness in proximal stimulus is a strong indicative of sameness in distal stimulus, due to the existence of normal environmental causes (Burge, 2010: 316-318, 325). Normal environmental causes, to which our perception is attunted, are crucial to overcoming sensory underdetermination. In fact, since perceptual constancies are directed towards specific environmental causes—color constancies, distance constancies, shape constancies, etc.—for perceptual constancy to be possible there must be a strong correlation between proximal stimuli and distal stimuli in the natural environment. So the existence of normal environmental causes is what makes perceptual constancy even *possible* in the first place.

Understanding how perception works—how organisms can perceive things as varied as colors and velocities—requires understanding what specific normal environmental causes such systems are attuned to. Attempts to understand perception as solving *general* constancy problems, rather than constancy problems based on normal environmental sensory data, generally fail to explain how our perceptual works. Many such models have failed to replicate perceptual constancy under naturalistic environments, as exemplified by the failure of two once prominent models of color constancy presented in Brown (2003: 257-9).

We have spent some time reviewing the fact that perception has to solve the underdetermination problem has necessitated perceptual systems to be organized in a specific way, a way that reflects environmental regularities pertaining to what such systems “aim” to detect. So much seems to be the object of contemporary scientific convergence.

We took the time to discuss this topic because it is exacly this apparent *purposiveness* in the organization of perception that has led Tyler Burge to attribute perception a representational teleology, and it is good that it is based on a characteristic of perception about which empirical researchers in general agree. In the second part of this dissertation, we will explain and evaluate the connection he sees between perceptual constancy, normal environmental causes, and representational teleology, and in doing so we will explicitly assume his interpretation of perceptual psychology.

We have already reviewed and dismissed four potential problems with Burge's understanding of perceptual constancy—consciousness-based accounts, relational accounts, the distality problem, and the lack of scientific convergence on the specifics—, and we will review how Burge's interpretation deals with four additional problems in the next section, including the notorious *disjunction* problem.

However, before we proceed to the next section, we wish to mention a criticism raised by Cohen (2015: §4) for the standard conception of perceptual constancy. He notices that our perceptual constancy abilities are not strictly speaking stable (*constant*), but are rather influenceable by contextual factors and often leads to error. Plus, our ability to detect constancy across different contexts is accompanied by an ability to detect *inconstancies* in these contexts. An object moving away at the same time appears constant in size and appears to shrink.

We have researched two empirical discoveries about perception that explain why perceptual constancy exhibits some instability, relative to the phenomena of contrast encoding and of perceptual adaptation. We bring sources not mentioned in Cohen (2015).

The first discovery is that contrast encoding is pervasive: the perceptual system has been shown to be systematically context-sensitive, in the sense that perceptual processing is mainly sensitive to *differences* of magnitude (for scalar attributes) or of magnitude vectors (for vector attributes), and uses this as a proxy to determine an absolute property (Cohen, 2015).

Here are a few examples. Our perception of an object's lightness and color depends on the surrounding lighting conditions (Cohen 2015: §2-3); and, in fact, retinal mechanisms for lightness and color detection function fundamentally in terms of contrasts (Pinker 1997: ch. 4, Stone 2011: ch. 3). More generally, vision science as a whole was reorganized some decades ago reflecting the widespread role of contrasts at all early stages of visual processing (Whittle 2003: Introduction).

Contrast encoding is paramout for many modalities of perception. The result is that "our perceptual reponses to an object/property are *not* constant, but instead change in interesting and systematic ways across variations in the perceptyal circumstances" (Cohen 2015: 628).

The second discovery is that perceptual adaptation plays a role in a number of perceptual processes. The perceptual system very quickly adapts to patterns of stimulation, in a way that significantly influences its interpretation of sensory stimuli afterwards (Webster 2003: 83-4).

Here are a couple of examples. Our perception of an object's taste is influenced by the previous state of our taste buds and, for humans lacking absolute pitch, tune perception is dependent on a comparison with recently-heard tunes (Cohen 2015: 626).

This effect is not limited to artificial laboratory conditions, as happens with strong color after-effects, but is in fact widespread in the way brains perceive the world and are involved in the maintenance of perceptual constancy. To pick two examples, light adaptation helps with lightness constancy and colour constancy (Webster 2003: 94-5) and contrast adaptation helps with contrast constancy, colour constancy, and constancies related to spatial properties (Webster 2003: 97, 101-2). Adaptation can lead our perception astray, since the way we perceive objects is influenced by our earlier percepts in a way that can mask their true properties.

However, even given the commonality of contrast encoding and perceptual adaptation, we do not think that this poses a challenge to the standard view of perceptual constancy. The perceptual variations we undergo are slight deviations from the true, invariant property being perceived.

As one researcher has pointed out, contrary to Cohen (2015), explanations in perceptual psychology involving perceptual constancies aimed at particular type attributes *do not* presuppose that perceptual constancies be absolutely stable:

> *"Individual's perceptions are* approximately *accurate with respect to* some *environmental psrticulars and attributes* enough *of the time to ground a form of explanation that takes states with veridicality conditions to be the product and participants in the law-like formation patterns being explained."* — Burge (2010: 88)

Furthermore, we have seen no researcher deny that perceptual constancy is not accompanied by perceptual *inconstancy*: our ability to detect the idyosincratic and perspectival differences across different perceptual contexts, such as the fact that distant objects *appear* smaller. The two abilities are simply co-occurent and this poses no challenge to the traditional understanding of perceptual constancy against which Cohen (2015: §2, §4) charges.

All in all, we have seemingly been able to fend off a number of potential objections to Burge's understanding of perceptual constancy and its relevance to perceptual science. There may be others which we have not come across and which might in fact be pointed out to us in response to this publication. Evaluating them would be the object of future investigation. We now proceed to examine the philosophical consequences of perceptual psychology as Burge interprets it.

Part Two

Naturalism and Reductionism

In this dissertation's first part, we have reviewed the interpretation offered in Burge (2010) of perceptual psychology. He has a specific view on the *typing system* employed by the science, that is, a specific view on how perceptual psychology carves up reality into explanatory kinds. His central claim on this issue is that it categorizes internal sensory states into *representational kinds*, characterized by their reference to distal type attributes. Furthermore, such attributes are exactly the objects of perceptual constancy mechanisms. Perceptual constancy is what institutes representationality in animals and human beings, under his analysis.

In the previous part, we also took the time to detail what perceptual constancy is, what roles it has played in the science, and how different scientists have come to see it. We have also seen how the thesis that perceptual constancy institutes perceptual reference can, indeed, solve the traditional philosophical worry about distal mental content.

In this part, we will explain in detail how Burge (2010) understands the notion of perceptual representation, with its associated notion of teleology, and the explanatory role this notion is supposed to play in perceptual psychology.

Discussing the nature of representation will lead us on to a controversy about *reductionism*, that is, about the possibility of understanding representation in terms of more basic kinds such as physical information, causal patterns (or "mechanical functions"), and biological functions, as has been proposed by various so-called "naturalistic" theories of perceptual content.

Despite his commitment to "naturalism" in the broader sense of constructing a worldview based on scientific theories and clear, rational philosophical theorization, rather than appealing to obscure notions contrived apriori, he rejects naturalistic theories of

perceptual content in the narrower sense presented above. He believes that representational teleology is irreducible to other kinds, despite figuring systematically in the successful science of perceptual psychology.

After reviewing prime examples of naturalistic theorization, to which Burge is opposed, we will rationally reconstruct his central objection to such theories, what can be called the *root mismatch objection*—that there is a fundamental incompatibility between representational teleology and the forms of teleology that can be grounded in informational, causal, and evolutionary patterns.

We finish by providing a critique of this argument under our reconstructed version. We argue that teleological explanations would be mysterious outside the scope of naturalistic theorization, and that fortunaately such theorization is compatible with perceptual psychology and the corpus of Burge's own theorization.

Section III – *Tyler Burge on the Role of Representations*

According to Burge (2010: 338ff), perceptual psychology takes perceptual representations to be teleological entities. That is to say, it takes them to be is something that bears a *purpose* of some kind. It can be evaluated relative to *success* or *failure*, depending on whether it has satisfied an associated *norm*. In his own words:

> *"Given that veridicality and non-veridicality cannot be reduced to success and failure (respectively) in fulfilling biological function, we must recognize a type of function that is not a biological function, a* representational function. *The basic representatonal functions concern representational success—veridicality, truth, making veridical, preserving truth, and so on."* — Burge (2010: 339)

In fact, he regards this as an *apriori* truth about our very concept of perception:

> *"Insofar as perceptual states have a representational function, perception is a teleological notion. And being inaccurate or non-veridical is a kind of failure or shortcoming for perception. Nothing could count as a perceptual state if it would not undergo a kind of failure if it were not approximately accurate. These are matters that we can know by reflection, without support from empirical considerations."* — Burge (2010: 535)

As stated, we will take his understanding of representation for granted. There is an ongoing dispute on what teleology characterizes perceptual states, as clearly not any teleology will do. Bodily organs may well have teleologies, such as the purpose of extracting macronutrients from food, but no reason has been found to endow our stomachs with perception. The teleology is not of the right kind. There is a specific kind of teleology which characterize what it is to be a perceptual representation; we may call it a *representational teleology*.

On Burge's understanding, a perceptual representation is a state whose teleology is to *have been caused in an appropriate way by an appropriate entity* (Burge, 2010: 345-6). More specifically, it is to have been caused *through perceptual constancy* by a token of its assigned type attribute. For example, a perceptual representation as of a *square* is just an internal state whose teleology is to have been caused through perceptual constancy mechanisms by a square token.

Perceptual constancy is itself viewed in normative terms. Given a teleological perspective, constancy mechanisms would not be those that just *happen* to group proximal stimuli in a way that matched their distal cause. They would be the ones with the *purpose* to do so (Burge, 2010: 408ff). As a consequence, a process which either momentarily or regularly did not match distal stimuli could *still* constitute a perceptual constancy mechanism, albeit one which produces illusion rather than veridical perception, for having failed its purpose. That is, in fact, how the very possibility of illusions would be explained.

Perceptual psychology categorizes internal sensory states relative to their representational teleology. That is, even if two internal sensory states have significant physiological and possibly functional differences, the science will treat them as naturally alike—as being of the same *kind*—just in case they bear the same representational teleology. In this sense, perceptual psychology employs a *normative* typing system. Teleological descriptions articulate the conditions for an entity's success or failure. Understanding semantic properties (such as representational content) in terms of normative conditions has been generically called *success semantics*.

Now, it is trivial to find examples of teleology in the social world, given that at least some minds can endow artifacts with purpose. Clocks undergo success or failure relative to their time-keeping performance. Bird nests are likewise evaluable with respect to stability and overall bird-egg protection (assuming that birds have the adequate intentional states). Dissertations can succeed or fail depending on their ability to adequately illustrate an

academic problem, depict key proposed solutions in the literature, and subject one thing or another to original scrutiny.

Perceptual representations are taken as teleological in the science for two reasons. The first reason has to do with the *explanatory target* of perceptual psychology, while the second reason has to do with its *explanatory format* (Burge, 2010: 37-40). We will take the time to explain each.

Onto the first reason: explanatory target. The science already assumes at the start, and aims to explain, the existence of veridical and nonveridical perceptual representation. As such, perceptual states are associated to a purpose, which can fail to be fulfilled, resulting in an unveridical state. The science organizes around explaining how our perceptual system is so vastly skilled in producing veridical representation, despite the informational difficulties associated to this, as reviewed in our discussion of sensory underdetermination.

Later, researchers started paying attention to how *unveridicality* (misdetection, misperception, illusion, hallucination) could shine light into how perception works. The reason is that there is significant structure behind *how* and *when* the perceptual system makes mistakes, in a way that can shed light into the specifics of perceptual processing and discriminate between competing models of perception. “One can easily imagine that the same degree of veridicality is achieved by different models. But a given pattern of errors is unique to a certain kind of visual software” (Gilchrist, 2003: 441).

For instance, our visual system would have the purpose of detecting external physical bodies. When we wander around our empty houses at night and become startled by what *seems* to be a shadowy figure, but is in fact empty dark space, we have undergone an illusion—our vision has failed to fulfill its purpose. Understanding what makes us misdetect bodies in badly-lit environments can shed light onto the specifics of our perceptual process, even if, on most cases, however, we are surprisingly good at discriminating objects and their exact countours even in noisy scenarios (such as rainforests and crowded spaces), a feat not closely matched by any visual detection computer program to date.

Either way, focusing on veridicality or unveridicality, the normative component finds itself present. It is in fact common sense that perceptual states can either be *veridical* or fail to be so. Early perceptual psychologists took this everyday distinction seriously when formulating questions, raising hypotheses, and constructing models. The distinction was cast in normative terms. The perceptual system was understood as having the of detecting

("matching", "mirroring") distal stimuli. Fulfilling that purpose constituted veridical perception; failing at that purpose amounted to illusions and hallucinations.

This general way to frame perception was the starting point for perceptual psychologists. Their questions and experiments were oriented towards explaining *how* veridical perception was possible and widespread—as it was squarely assumed that it *was*—, and *why* it sometimes failed and ended up producing illusions, as retold by Gilchrist (2003: 438-41). The ability to produce veridical perception still guides large swaths of contemporary scientific research in perception (Burge 2010: 310-2). This is so because, as it turns out, questions, hypotheses, and models guided by a normative perspective proved themselves to be much more fruitful than competitors. The normative representational framework truly helps in framing problems and solutions.

Onto the second reason why perceptual representations are understood as teleological: explanatory format. Teleological descriptions are appealing because the provide a distinctive explanatory form: teleological explanation. Perceptual representations are also posited in perceptual psychology to take part in explanations, and these explanations take the form of a teleological explanation. These are explanations that appeal to purposes in explaining why something happens.

Teleological explanations aim to explain what entities *are* based on what they are *supposed* to do. Teleological explanations are part of everyday life for teleologies instituted by minds. We can probe more effectively, and also explain more about, the organization of a clock by framing clocks as devices designed for time-keeping. The idea is that, given such a teleology, the clock is expected to be organized in such a way as to reliably *fulfill* that teleology.

Applied to perceptual representation, or some other naturally instituted teleology, the idea is that we can better probe and better explain the perceptual system by framing perception as organized to fulfill a certain task. For instance, consider the assumption that the early visual system purports to detect oriented lines in the environment. One would then *expect* that this system would be configured so as to reliably detect oriented lines. This helps scientists to probe the visual system more effectively as well as to provide systematic explanations that make much more sense, because these explanations would be unified by the perceptual system's central purpose: the fulfillment of its teleology, which we call *veridical representation*, and which consists in having been caused by a certain type attribute through a constancy mechanism. Regarding systems in which various initial points robustly converge to

the same result, some in the literature have claimed that goal-oriented language can provide a fruitful perspective (see the remarks on Braithwaite and Nagel in Wright 1976: 28-9).

It is illuminating to group together natural processes (such as sensory processing mechanisms) based on the problem they *aim* to solve, even if their solutions are quite different and regardless of how physically distinct they happen to be. Bearing this information helps us articulate how the systems work as well as to predict their functioning.

Now, mind-instituted teleologies, such as a clock's, seem easy enough to understand so long as we take for granted the previous exisence of minds. What makes teleology a thorny topic in philosophy is that researchers sometimes have reason to believe that there has been a teleology instituted naturally—mechanically, mind-independently—that is, without being grounded in an antecedent teleological entity, with minds being the prime example.

That means that objects acquire norms without there having been any agent to select it, design it, or somehow else attribute a purpose to it. There is little consensus on how (and even *whether*) such a *natural* teleology could come into existence.

This raises the question: does perceptual psychology have any light to shed on how perceptual teleology is instituted? There seem to be two possibilities.

The first is *non-reductionist*: perceptual psychology assumes, as a working hypothesis, that perceptual types are characterized by a representational teleology—they have success conditions wherein success is veridical perception and failure is perceptual illusion—, and that hypothesis has proven itself necessary for a fruitful science. This approach offers no method, no principled way to attribute teleology, and is in this sense a haphazard, a *hit-or-miss* approach: perceptual states are simply assumed to have basic teleologies and that assumption is preserved each time just in case it proves both fruitful and ineliminable.

The second is *reductionist*: the perceptual system, along with its phylogenetic and ontogenetic relation to the environment, can be fully described under a *non-teleological* language, but which can be systematically related to a teleological description. That is, perceptual kinds can be described non-teleologically, rather than being fundamentally characterized as representational kinds. This would be the approach taken by naturalistic theories of mind in the aforementioned sense. We will discuss reductionist proposals in depth in section IV.

The above is the case with clocks and other artifacts: a full mechanical description can be systematically related to its purpose to keep time. The design-oriented teleological description would be vital for researchers grasping for clues on how clocks work, and it

would be even more vital if clocks were as complex as brains, for teleological descriptions provide a unifying thread from which to interpret the whole system.

Burge (2010) can be unclear on which approach it prefers. On numerous occasions, an internal sensory state's *causal role* is offered in the book as necessary and sufficient reason to attribute it a reprensetational teleology. Every entity that has historically exhibited, or which presently exhibits, or which is disposed to exhibit the right causal role—name, that of being produced with law-like regularity by a constancy mechanism—is counted as a perceptual representation. Representational content is determined by what environmental occurrence that sensory state is causally associated through perceptual constancy with law-like regularity. The following quote is representative:

> *"What are constitutive to perceptual-state kinds are the causal relations to environ-mental attributes, and the individuation of laws of formation of instances of perceptual-state kinds by deep statistical or law-like regularities in the physical environment that are specifically relevant to attributes being represented. The perceptual state as of light-ness depends constitutively for being what it is on causal relations to lightness, or to attributes systematically related to lightness."*
> — (Burge, 2010: 365)

However, the book also makes it perfectly clear that there are no grounds to accept reductionism about perceptual states; this is repeatedly emphasized by Burge (2010). That is, the book understands that, fundamentally, teleologies are attributed non-reductively, as per the former, unprincipled approach mentioned above. He repeatedly states, in different detail and highlighting different reasons, that reductionism *could* be true, but that there is no reason to think that it is, and although (as we shall see in Section V) there are some reduction proposals he believes are *apriori* determinable to be false:

> *"We have no reductive explanation for what it is to be a perception as of something's being spherical, or for what it is to represent something in an objective way. I do not expect such explanations. I will initially assume and later argue that, as far as we can now tell, psychological kinds are explanatorily primitive, in the sense that specifications of them are not exhaustively reducible in scientific or other explanatory enterprises to specifications that are not distinctive of psychology."*
> — Burge (2010: 58)

Despite there being no reduction available, he believes that there is no *apriori* reason to believe that there will always be a reduction avaiable, so that we can accept non-reducibility when it has proved itself essential to constructing a successful science. These two excerpts point this out with distinctive eloquence:

> *"I think that explanations in representational terms do not reduce to explanations in other terms: psychological states described in such terms are explanatorily indispensable. Some of the most rigorous, powerful parts of psychology use representational terms. Explanation of perceptual accuracy and illusion, and explanation of the formation of perceptual states, are ineliminably in representational terms. Such terms are a secure part of science."*
> — Burge (2010: 81-2)

> *"An antecedent commitment to reduction is ungrounded ideology, not an expression of science or reason. (...) Reductions are a legitimate type of explanatory unification. Occasionally reductions succeed. In principle, representation might be somehow reducible to other notions. I believe, however, that trying to reduce representation and veridicality to something more "naturalistically acceptable" is probably pointless and hopeless. At any rate, the reductionistic proposals that have been made so far seem to me hopeless.*
>
> *Notions like representation earn their keep in science, and to a large extent in common sense, by figuring in successful explanation. Successful explanation is marked in the usual ways by yielding agreement, opening new questions, making questions testable and precise, engendering progressive improvement in theory and experimentation. Mainstream work in perceptual psychology displays these features."*
> — Burge (2010: 298)

Primitive terms are those which do not receive explicit characterizations, that is, definitions using other terms. Although 'reduction' is another polysemic term, one sense in which it is used is to designate an explicit definition. For instance, should Humean empiricists be right that causal relations can be defined by counterfactual relations, then causation would be *reduced* to modality. Reductions are never intended as mere nominal exercises—as when one says "henceforth by 'causation' we mean so-and-so"—but rather as substantial claims of some sort of *identity* between two features of reality which have been independently discerned. Other theorists claim that causation cannot be reduced to anything else and, as a result, has to be accepted as primitive.

Lacking a reduction, a primitive term's meaning is either taken to be *intuitive*—whose meaning and application conditions we can discern despite our inability to articulate it—or to

be *implicitly defined* by their usage—whose meaning and application conditions are quite precisely constrained by the role it must play inside the theory. *Origins of Objectivity* argues that a notion can be accepted as primitive when no reduction is available (despite best efforts) and, yet, theorists have relied on that notion to construct their successful theories. While some researchers claim this to be the case for causation, the book argues it to be the case for teleology in perception. This topic raises heated debate in the philosophy of mind community as most theorists believe that everything should be reducible to descriptive facts, i.e., that normativity cannot be taken as primitive on pain of obscurity and mysticism.

As we have seen, Burge begs to differ, and is resolutely committed that representational teleology might well be irreducible. However, what then should we make of Burge's commitment to law-like causal regularities being constitutive to individuating perceptual representation, given that representational kinds are irreducible to causal regularities?

There is a way to render the two commitments compatible, although this is not made explicit in the book. It is an original contribution of ours. What could be said is that perceptual psychology performs a *gambit* when it states that (*i*) a sensory state is embedded in the right law-like causal regularity if and only if (*ii*) a perceptual representation associated with a certain teleology. What has happened is that, *as it turns out*, the gambit has recurrently proven successful. The science had no guarantee that this would be the case, and has no guarantee that it will continue to be so. If the science had purported to *reduce* teleology to causal role—if it had aimed to define the former in terms of the latter, or otherwise apriori deduce the former through the latter—this would not have been a hit-or-miss gambit, but rather a sure hit.

If Burge is correct in his repeated emphasis that there is a fundamental distiction between representations and causal roles—and anything which can be derived thereof—then perceptual psychology would have to countenance the possibility of '*i*' occuring without '*ii*' and vice-versa. We return to this topic in detail in section V when we discuss why Burge believes the two to be fundamentally distinct.

To sum up, there has been so far a covarying relationship between causal role and teleology, but it is not articulated in a principled way: one cannot derive teleologies from causal roles. Rather, each teleology is posited anew and is found to work after due empirical work. This sits Burge's *Origins of Objectivity* (2010) squarely within the non-reducionist camp, despite his admission of the aforementioned covariance, which is characteristic of a reductionist approach.

A central point to understand is that not every teleological description is interesting. Some teleological descriptions can be replaced by non-teleological ones without hindering our ability to understand a domain. In that sense, their teleological aspect was *eliminable*, and that fact raises the suspicion that teleology was but a descriptive gloss, merely one way to frame the situation. What draws the scientist's attention, says Burge (2010), is an *ineliminable* teleological description—one whose explanatory power cannot be found elsewhere. This is because it is only successful and ineliminable teleological descriptions that make it plausible that the world objectively contains teleological aspects.

We claim that there can be ineliminable teleological descriptions even if one adopts a reductionist stance towards teleology. That is, teleological descriptions can provide unique explanatory gains even if teleological descriptions are reducible to non-teleological descriptions. We will defend this claim shortly. We now turn to discuss what it is for a teleological description to be ineliminable.

As we have seen, Burge (2010) makes an extended argument that perceptual constancies are what render ineliminable a representational (and thus teleological) account of perception (cf. Burge, 2010: 63n3). We will now be in a position to understand why. Whenever perceptual constancies are present, researchers find it expedient to provide representational explanations. It is only when perceptual constancies are missing that perceptual psychologists, at least over time, find a way to provide a non-teleological explanation *without explanatory loss* in lieu of the original representational explanation.

Here is an instructive, detailed example of an *eliminable* teleological description. Valentino Braitenberg's *Vehicles* (1984) exemplifies many machines that could be described teleologically, but each in an eliminable way. We ask the reader to grab a pen and paper and to follow through the simplest such vehicle.

It is a small, brick-shaped contraption with only four relevant features: two motorized back wheels (right and left) and two front light sensors (right and left). The wheels are independent and turnable, and the vehicle does not topple. The vehicle has some internal wiring. The right sensor connects to the left wheel's motor, the left sensor connects to the right wheel's motor; that is, the sensor-wheel connections are crisscrossed. Whenever the *right* sensor is activated, the *left* motor turns the left wheel forward, and vice-versa.

Things being so, suppose that light is *uniformly* shined upon the vehicle's front face. Both sensors will be equally activated and, therefore, both motors will be equally stimulated and, as a result, the wheels rotate synchronously. The result: the cart moves forward in a

straight line. Now, if light were to be shone more strongly upon the vehicle's *left* side, the left sensor would activate more strongly. Inevitably, the right wheel would rotate faster, so that the cart would start making a *leftward* turn, effectively going towards the light source. (The reason behind the leftward turn can be a bit difficult to visualize.) *Mutatis mutandis*, the same holds for the right side.

As a result of our simple setup, we have a cart that follows light sources whether they come from the left, the front, or the right. Here comes the interesting bit. Note that we *could*, if we wanted to, describe the cart as bearing the *teleology* to follow the light. After all, in some watered-down sense it does seem to "want" to go towards the light source, and that's what it always does. We may then try to understand the cart under this framing, providing teleological explanations. Doing so is easy and cognitively comfortable for human beings; for many people this would be the first description they would come up with, if they did not know the vehicle's internal wiring.

However, it is perfectly possible to describe the cart in purely mechanical terms, citing only sensorial stimuli and internal signal processing. Such a mechanical description is, in fact, exactly what we provided above when outlining the cart's setup. This means that our proposed teleological description is *eliminable*.

Whenever a similar situation happens in perceptual psychology, Burge claims that the science drops the eliminable teleological description. That is, whenever a creature's perceptual system (and perceptual-guided behavior) is comprehensible in mechanical terms, mentioning only proximal stimuli and internal signal processing, over time the easy and comfortable teleological descriptions are abandoned in favor of informational-mechanical ones. As shown by many examples in Burge (2010: Ch. 10), this has indeed been the fate of many sensory processing mechanisms.

Systematic understanding them through law-like generalizations does not require teleological language. The brain's processing can be described using notions that make some conceptual sense, involving signal-processing mechanisms such as *comparison*, *summation*, *averaging*, *weighting*, *accumulating*, *habituation*, *conditioning*, *serial sampling*, *pattern matching*, and so on, which are operating entirely over sensory data such as the *location* and *intensity* of type stimulations. Such mechanisms are thus said to be *merely* sensory.[5]

This is not possible when perceptual constancy mechanisms are involved. Perceptual constancy filters proximal data which reflect the organism's perspective and bodily

5 See the section Examples of the Sensory-Registration / Perception Distinction in Burge (2010: 421ff).

idiosyncrasies from proximal data which reflect external invariant properties. To do so, it must ignore a lot of proximal stimuli and bring in a lot of implicit assumptions about the environment. An approach based on signal processing is, as a result, uninformative and in fact misleading. Such an approach could not treat radically distinct stimuli in astonishingly similar ways. The signal processing cannot be described as processes of comparison, averaging, habituation, serial sampling, and so on; it cannot be made sense of using such notions. Similar stimuli are mysteriously treated differently; very different stimuli are mysteriously treated alike. The signal processing that is going on can only be made sense of, and its mysteriousness dispelled, in light of its *aim*: to discriminate an invariant external attribute.

The upshot is that *perception* proper, and thus perceptual *representation*, enters scientific theorization only when a teleological, representational framework proves itself explanatorily ineliminable. As it turns out, this occurs thanks to perceptual constancy.

An explanatory gain is important to sort out what is truly representational and what is not. Thermometers can be described in representational terms, but doing so offers no explanatory gain. Some sensory mechanisms share that fate with thermometers. For others, a representational outlook provides distinctive insight, and that is when scientists become confident to posit genuine teleology.

Enough has been explained about the proposal in Burge (2010) to address some challenges that have been raised in the literature for theories of perceptual content. We believe that Burge's teleological view of perceptual constancy is sufficient to address three concerns that have been raised by Schulte (2021: §2.2) relative to the notion of perceptual constancy. We also believe that such a teleological view, added to his considerations about normal environmental causes, can also solve the *disjunction* problem, also called the *indeterminacy* problem (Fodor 1990: §2, Schulte 2018: §5, Roche & Sober 2019: §2.1).

Schulte's worries are directed toward the notion of perceptual constancy characterized as: the process of forming the same perceptual state in response to the same distal environmental condition, despite variations in proximal stimuli—but with no mention of teleology.

The first concern he raises is that there is a deep unclarity about which variations in proximal stimuli are taken into account in the above definition. His claim is supported by a single example, namely, that carnivorous plants clamp down if and only if at least two of their trigger hairs have been touched in a small interval of time, *ignoring* both the precise timing and the total amount of hairs. Such clamping down is biologically useful because usually only

edible objects trigger two or more hairs in such a short interval. Schulte's point is that, if variations in proximal stimuli are being ignored, then per the above definition it would follow that carnivorous plants exhibit perceptual constancy relative to edible objects—an implausible conclusion.

The solution to this conundrum is that the carnivorous plant's functioning can be described without appealing to a teleology to detect edible objects. Its mechanism is very simple, as described above. Perceptual constancies are introduced by scientists only when a sensory mechanism cannot be interestingly described *except* with an appeal to an overarching *purpose*, which stands as the sole organizing principle behind an otherwise inscrutable sensory apparatus.

The second concern he raises is that there is no standard criteria to attribute perceptual constancies. This is very much true, and it is true as well for representaitonal teleologies. What is not true is that this is worrisome, because perceptual constancies and representational teleologies are introduced when they have proven useful in scientific theorization. They have "earned their keep", as Burge formulates it.

The above consideration also solves the third concern he raises about the standard account of perceptual constancy: that it is neutral relative to the *mechanism* that implements it, with the negative entailment that quite non-constancy-like mechanisms would count as perceptual constancies. The example he provides is: ignoring features of the proximal stimuli which mask the distal cause—the necessary and sufficient condition for perceptual constancy, on the account he critiques—could, in principle, be performed by a look-up table that exhaustively and directly connects all possible inputs to the correct perceptual state.

The solution is that such a look-up table could be described non-teleologically, and therefore on a teleological account of perceptual constancy it would not count as perceptual constancy. We could also add that Schulte is failing to imagine the immensity of this look-up table, given that perceptual constancy mechanisms take a myriad of contextual details into account, such that the same retinal stimulus could be interpreted differently under different conditions. We should always be wary of artificial examples that ignore the complexity of the real world. But this is just an aside.

So much for Schulte's worries. Now, we have mentioned that considerations relative to normal environmental causes can also solve the *disjunction* problem. This problem stems from the fact that any object can be specified under multiple different descriptors. For instance, a mosquito may be described as a black surface, as a nutritious blob, as a small

flying insect, or as a mosquito. If a perceptual system exhibits perceptual constancy relative to mosquitoes, the challenge is to specify what determines whether it is a perceptual constancy relative to black dots, relative to nutritious edibles, relative to mosquitoes, or relative to some other descriptor. The disjunction problem is the challenge of answering: is there one privileged and precise type attribute to which the animal exhibits perceptual constancy and, if so, why?

The solution can be gleaned from the fact that perceptual psychology is not an isolated science, but one that aims to integrate with a broader range of biological sciences, as we have pointed out in this dissertation's introduction. Explanations in perceptual psychology must connect with explanations in such biological sciences. As such, and Burge makes this very clear, perceptual constancies are envisioned as matching *biologically relevant kinds*. Among all possible descriptors of an object that could be classified as the "normal environmental cause" of a percept, researchers in perceptual psychology single out descriptors that make sense from a biological point of view and that systematically integrate with the others descriptors standardly employed in fields as diverse as evolutionary science, cognitive ethology, and ecological science. As per Burge:

> *"The formation principles tend to serve the representational function of the perceptual system in providing veridical perception of entities in the environment. The relevant entities are the explanatorily relevant environmental antecedents of the proximal light arrays. The theory assumes that perception represents elements in the distal environment. This intuitive assumption is grounded in a larger explanatory point of view. What count as potential perceptual objects—as relevant distal antecedents—are roughly those that can be discriminated under certain conditions, that the internal processes are best explained as bearing perceptual constancies with respect to, and that are ecologically relevant to the individual's basic functions—functions such as eating, navigating, mating, fleeing danger."* — Burge (2010: 94)

> *"The fact that biological functions of sensory systems are relatively close to representational functions makes psychology possible. The fact that biolofical functions are not the same as representational functions helps make psychology independently interesting."*
> — Burge (2010: 303)

To continue with our example relative to mosquitoes, to resolve that dispute it must be pondered what descriptor the other biological sciences would employ when describing the perceiving organism's environmental interactions. Were the organism to be a frog, would

biologists understand their ecologically relevant interactions to include interactions with *black surfaces*, or with *flying blobs*, or with *small flying insects*? Most likely, researchers would privilege the descriptor *mosquito*, corresponding to the ontogenetically and phylogenetically relevant species with which the frog interacts.

It is by perfoming considerations such as the above, Burge urges, that researchers in perceptual psychology can to a large extent solve the disjunction problem, since alternative descriptors are largely incompatible with the relevant associated sciences. This is not to say that perceptual reference to type attributes is exactly determined, for in fact perception is best conceived as performing approximations (Burge 2010: 379-80, 388), but the range of underdetermination is vastly narrower than the disjunction problem would originally make it seem.

Having seen how Burge's teleological account can solve standard worries about perception, it is time for another round of clarification. We have not, so far, made it in any way clear what a teleological explanation is, that is, what sort of explanatory gain is provided by an ineliminable teleological description? We now turn to fill this gap. Doing so will allow us to better understand what teleologies are even supposed to be in the first place, a discussion that will be central in the following pages in this dissertation.

We uphold that there is not a single passage in *Origins of Objectivity* (Burge 2010) where the nature of teleological explanations is laid out. As an original contribution, we attempt to lay it out. *What* do they explain, and *how* do they explain it? We found only one way to make sense of teleological explanation, other than claiming that it is a 'fundamental explanatory form' that can only be understood on its own terms—which seems like an obscurantist sleight of hand.

Suppose that *A* has the teleology to do *B*. That is a teleological fact, one which would be presented in a teleological description of a science involving *A*'s and *B*'s. The idea is the very fact that *A* has the purpose to do *B* helps explain why *A* reliably does *B*. Teleologies explain why they themselves are reliably fulfilled. That is *what* teleological explanations explain.

Teleologies explain their own fulfillment; that is sure a strange proposal, for it is circular in some sense. We reached this realization when we noticed that Burge's understanding of representational function is, in the very same sense, circular (Burge, 2010: 309-10).

He claims that representational teleologies ground the explanation of why we have the ability to reliably form veridical representations—that we have such ability, he claims, is something perceptual psychologists assume and have as an explanatory target, as we reviewed in section II. However, he defines veridical representation in the same as we did, namely, as «the fulfillment of the representation's teleology». Representational teleologies thus ground the explanation of why they themselves are reliably fulfilled. That would be like explaining *why* the heart pumps blood by noting that the heart has the *teleology* to pump blood. In the above scheme, hearts are *A*'s and pumping blood is a *B*. "Visual representations as of *squareness*" (an A) help explain why "the visual system is able to reliably detect squareness through perceptual constancy" (a B). Burge was seeming to have made an unnoticed mistake.

Then it dawned on me: that is exactly what Burge is asserting. It is no mistake. Teleologies, at least in the natural world, explain their own realiable fulfillment. He is not explicit about it, as we have said, but we have found no other way to make sense of his claims. This is different from what other prominent theorists regard as the target of teleological explanations. Garson (2016: 36-7) understands that, if it is a fact that *A* has the teleology to do *B*, then this fact explains why *A* exists (in order to do *B*). Perhaps so. But we take that fact *also* to explain why *A* reliably does *B*.

But if that is *what* teleologies explain, then *how* do they explain that? Philosophical puzzlement with teleology has a clear and definite source, at least among those with a naturalistic perspective on the natural world: the assumption that teleological properties have no causal properties—at least when causality is understood narrowly, as involving some sort of *mechanism*. An object's *purpose* will perform no causal work; teleologies are not intended as just another causal account (Wright 1976: 25-6). Purposes are causally efficacious only when they are the purposes of a mind, which is not the case here. We concur with this naturalistic assumption, and nothing in Burge's work indicates that he believes the contrary or, more importantly, that he is able to make sense of the *negation* of this assumption.

Teleologies, thus, cannot cause their fulfillment. The fact that one has or exercises the capacity to represent *distance* does not *cause* one to be able to do so, for example by causing one's perceptual system to contain distance constancy mechanisms. Teleologies cannot provide causal explanations. Teleological explanation, therefore, must be of a different kind.

As far as we can see, this question is not clarified in Burge (2010). We have given this question a great deal of thought and we have come up with a thesis that answers this question. We have found no other way to make sense of teleologies explaining their own fulfillment.

While we are unable to argue that there is no other solution, we assume tentatively that there is no other one because we have not come up with any, and we explore the consequences of that assumption.

Our idea is that there is a constrait on the institution of teleologies. That constraint would have to be satisfied by any theory of *foundational normativity*, that is, by any theory about how teleologies come to be instituted in the natural world, if teleological explanations are to make sense.

The constraint is that it would only be possible for an object to institute a teleology in case that very same object had, at the moment of institution, a mechanism that reliably fulfilled that teleology. If we are correct, then teleologies are *constitutively*—that is, by their very nature—strongly correlated to their own fulfillment, for at the least at the moment of inception there was a reliable mechanism set into place. The system had all its internal components organized so as to guarantee such reliable fulfillment.

This explains why teleological descriptions can be *illuminating*. Given that constitutive correlation, they provide *information* about the way a system is organized, because it must at least once have been systematically organized to fullfil their teleologies. Teleologies must be constitutively associated with their *explanandum*.

Ultimately, the *causal* explanation for *why* the system is the way it is does not ever reside in teleology. One can never answer causal why-questions about perception using teleological descriptions. This has led some researchers, which do accept that teleological language is meaningful, to deny that there can be any teleological explanation at all (cf. examples in Wright 1976: 99-100). Yet, scientific why-questions extend beyond queries for causal explanations. Teleological explanations provide non-causal explanations.

Such explanation has the exact shape described above: since everything in a teleological system is oriented towards fulfilling a teleology, teleological descriptions constitute a framework which helps us to understand *how* a system is internally set-up and to predict *what* it will do. If we ask, “why is the system like so-and-so?”, in an extended sense we can answer that it is *because* being like so-and-so is expedient to fulfill its teleology. But teleological explanations do not answer *why*-questions in the stricter, causal sense. Teleologies merely *lawfully correlate* with the correct causal explanation. They are informative about causal why-questions without directly answering them.

Here is another way to see how teleological explanations can be informative in perceptual psychology. Perceptual constancy relative to some external attribute *A* requires

ignoring many details in proximal stimuli, by filtering what is idiosyncratic to one's body and perspective and what, in fact, carries information about the invariant external property *A*. As such, detecting *A* serves as an end-point towards which the system robustly progresses to when starting from radically distinct proximal stimuli. Understanding perception in terms of such end-points and the conditions in which they are reached (veridicality) or not (illusion, hallucination) is highly useful.

Recall, however, that teleological explanations must be *ineliminable* for them to be acceptable in science. What this means is that they must provide an explanatory gain not otherwise attainable. This explanatory gain, we submit, to provide an organizing principle.

Teleologies serve as an organizing principle in the following two senses. First, some causal mechanism are so intricate that there is no informative, systematic way to describe its operations *except* by casting the mechanism as containing steps torwards fulfilling an overarching purpose. Second, a teleological perspective provides meaning to the notion of *malfunctioning*, which is often useful when observing natural mechanisms. Many scientific fields, ranging from ethology to psychiatry, rely on the fact that there is something which a mechanism is supposed to do, but sometimes does not. Those are the explanatory gains provided by teleological explanations.

We have just provided a constrait for theories of foundational normativity that wish to make sense of teleological explanation. Teleologies are objective norms constitutively associated with their *explanandum*—namely, a reliable performance that satisfies the teleology—and which allow us to understand otherwise incomprehensible mechanisms.

We can now inquire on what process could institute teleologies could satisfy this constraint, that is, what theories of foundational normativity would be adequate relative to it.

We have again only come up with a single answer: a theory in which teleologies are instituted by a *causal selection process*. Entities that were mechanically selected due to their reliable causal effects—that is, selected through a causal selection process—would be thereby attributed the teleology to exert those causal effects. This makes clear whence comes the constitutive connection between teleologies and their reliable fulfillment. A teleology to do *x* is only instituted when some mechanism reliably did *x* and, as a result, was mechanically selected. The beauty of it is that the very same selection process which institutes the teleology to do *x* also *propagates the mechanism which reliably does x*.

Selection processes *causally* explain why subsystems have teleologies and meet the conditions for reliably fulfilling it. A selection process can at once causally explain why a

visual cortex segment has the teleology to detect shapes and why it is competent at such a task. As teleologies are causally inert, they are uninteresting except for their constitutive relation to selection and, indirectly, to *actual function*—to effective causal role—which allows it to guide research and serve as central organizing principles of our models.

Should we be right, representations could play no explanatory role in perceptual psychology unless their distinctive *representational teleology* had been instituted by a causal selection process. Otherwise, they would be unable to explain *any* aspect of our perceptual apparatus, as there would be little reliable correlation between teleologies and their fulfillment—and, thus, no guarantee that the perceptual system would be organized so as to fulfill its teleofunctions. Teleologies would then neither illuminate nor explain anything: neither the existence of perceptual constancies, nor our reliable capacity to perceive veridically.

Under this view, teleologies end up as nothing but a historical property: to bear the natural teleology to do *x* is, precisely, to have been mechanically selected to do *x*. To the mechanistically-minded philosopher, teleological properties could not be a present-tensed property, for such a philosopher cannot countenace a *causally inert* property, and teleologies indeed have no present causal powers. When cast as a past-tensed property, it can be characterized in terms of *past* causal roles. In this case, having gone through a causal selection mechanism.

Should the above notion of teleology and teleological explanation strike the reader as excessively watered down, bearing little resemblance to what these things are generally thought to be, then we have just incurred in a long-standing pattern of deflating teleology (Wright 1976: 3-4). We believe this to have been inevitable.

In the next section, we will review a class of theories about foundational normativity and perceptual representation that fall under the category of *teleofunctionalism*. The considerations we have made in the preceding pages are congenial to teleofunctionalism. These theories are also often submitted to the charge of excessive deflationism, to the point of claiming that teleological descriptions are merely ways of speaking. But, as we shall see in section V, teleofunctionalism is at serious odds with considerations put forth in Burge (2010).

Section IV – *The Philosophy of Teleofunctionalism – Literature Review*

We have seen that *perceptual constancy mechanisms* play a central role in perceptual psychology. Following Burge's interpretation of that fact, we have reviewed how perceptual constancies guide scientists in fruitfully positing *representational teleologies*. This raised two questions. First, what is the relation between teleological properties and the remainder natural order? Second, what makes a teleological explanation even possible?

In this section, we will explore a family of positions known as *teleofunctionalism*, to which our theses worked in section III seem to belong. Since our aim is to better understand what Burge is opposed to, and therefore better understand his views, it would be outside the scope of this dissertation to perform a thorough literature review. As such, we will not discuss in detail all central historical authors in the teleofunctionalist camp; doing so that would be more fitting for an academic work directed at teleofunctionalism, rather than at Burge and perceptual psychology.

Furthermore, we will not attempt to defend one version of teleofunctionalism over others. Although we do have our preferred view, which we have already stated and will restate below, Burge's criticism to teleofunctionalism applies to all major versions. All of them are related to causal effects and biological fitness, which Burge views as fundamentally incompatible with the nature of representational teleology.

Having said so, let us turn to the two questions above. On the first question, in section III we have mentioned Burge's emphasis that there are law-like regularities connecting environmental attributes, sensory processing, and representational teleologies. Similar remarks have prompted many researchers to propose that representational teleologies can be systematically related, and in this sense *reduced*, to non-teleological historical and causal patterns in nature. As we'll see in section V, he argues that evolution- and causation-based reductions are impossible, although others may be possible (Burge 2010: 298).

On the second question, no discernible answer seems to have been even as much as suggested in Burge (2010). In section III, we laid out the only attempt at explicating *teleological explanation* that we have found reasonable. Our explication, however, puts us at odds with Burge's contention that no causal selection mechanism could institute representational teleology. What we have offered places us in the teleofunctionalism camp, to which Burge vehemently opposes.

Since teleofunctionalism rests among the major positions in the debate about perceptual and mental content, and since it is repeatedly denounced by Burge as inadequate—incompatible with perceptual psychology and, in fact, with the very notion of perceptual representation—, we have found it expedient to review the major teleofunctionalist proposals that have been offered. We will relate the position we have committed to—based on our explication of teleological explanation—to some major available positions in the teleofuctionalist camp. This will allow us to better discern what Burge is criticizing and, thus, in section V, to better evaluate whether his criticism is sound.

Let us begin. The philosophical background in which teleofuctionalism is developed involves at least two layers of problems. The first layer comprises of problems relating to mental content. The problems most prominent in this literature are about representational *reference* to particulars, token attributes, and type attributes. We have already seen the distality problem—regarding the perception of the distal rather than proximal causes of percepts—as well as the disjunction problem—on the multiple possible descriptors for any particular object of perception. We will also discuss the *misrepresentation* problem—about the possibility of representing what is not there. We will not be concerned with problems about representing abstract and fictional objects. Neither will we take time discussing what is called *cognitive content*, which includes perspectival aspects of perceptual representation (discussed in Burge 2010: 391-2) and the modes of presentation under which objects can be represented.

The second layer comprises of problems relating to teleology: what it is, how it could be instituted (that is, come into existence), and its explanatory role. A major problem in this domain is determining whether teleologies are determined *narrowly* or *broadly*. Teleology is narrow just in case two physically identical objects would necessarily be subjected to the same teleologies. Teleology would be broad in case environmental and/or historical factors were necessary to determine the norms associated to an entity. Some theorists also believe that there are meaningful notions of both narrow and broad teleologies. (See Brown 2006: §1 for more details.)

The above problem will not figure prominently in our discussion. The reason is that Burge and teleofunctionalists concur that representational teleologies, and thus perceptual content, are determined broadly. Perhaps both can accept there to be a form of narrow content, but their chief concern is with broad content. Burge has offered no systematic theory about how teleologies are instituted, but he emphasizes that they can only serve an explanatory role

in science because of their systematic connections to environmental and historical factors (Burge 2010: 94, 365, both quoted above). We will, however, mention the ways in which different theories understand perceptual teleology to be broad.

Other problem sets can be brought to bear on this issue, with epistemic problems being a prominent example; as laid out in this dissertation's preface, many problems are intertwined. Rather than broadening our scope to include all relevant discussions, we will in this section focus on problems pertaining to teleology: its nature, institution, and explanatory role. Teleology will be our guiding exploratory thread, with mental content problems being discussed on the way.

Part of the dispute over perceptual teleology stems from differing views on how natural teleologies could be instituted—that is, how something could come to *acquire* a teleology without any minds around. Let us refer to such views as *foundational* theories of normativity. Given our focus in teleofunctionalism, we set aside researchers who dispute that perception contains any teleology at all.

Telofunctionalists disagree about how natural teleology comes into existence, which leads them to diverge over what teleologies our perceptual apparatus *could* have even in principle, let alone over the teleologies such apparatus in fact has. There is, however, a large arena of agreement between them, which we will now review.

There are varied kinds of teleology. Perception is thought by teleofunctionalists to bear a distinctive form of teleology. Teleofunctionalism proposes that such distinctive kind is a *proper function*. Perceptual content consists in proper functions associated to the perceptual system—that is the broadest characterization of teleofunctionalism about perception. We will see how this is further specified in the views of specific teleofunctionalists.

Teleologies, broadly conceived, are any norm associated with an entity. Any aspect of an entity can be subjected to a norm. A teleology becomes a proper function just in case it is a teleology that evaluates an entity's *causal relations* ("causal role", "functional role"), that is, the way in which it interacts with other entities. If we call *mere* functions that which entities in fact *do*, in a causal sense, then *proper* functions are what entities *should* do: what mere functions they are supposed to exhibit. This terminology itself is arbitrary and was originally introduced by Ruth Millikan (as retold in Millikan 2002: 116), but the distinction between mere ('accidental') function and proper ('purposive') function is crucial to the concept and should be preserved in any explication (Millikan 2002: 117; Garson 2016: ch. 1.2; Allen & Neal 2020: §2).

The three questions we have listed are answered differently by teleofunctionalist theorists. First, what is the nature of the teleological aspect of a proper function? Second, how can proper functions be instituted? Third, what can proper functions explain, and how? To better understand teleofunctionalism as a research programme, we will look into theories of proper functions.

Proper functions occur in a wide variety of scientific fields (see examples in Garson 2016: ch. 1.1), which caused the notion to receive a lot of philosophical attention. This discussion should be guided by scientific research: our resulting explication of what proper functions should allow them to play the explanatory roles they have been designated in scientific practice (Garson 2016: ch. 1.2). We must also be open to the possibility of *pluralism* about proper functions; there might be more than one notion being employed in the sciences, given the world's complexity and the fact that the notion is used across fields as diverse as evolutionary biology, ethology, ecology, cognitive science, physiological medicine, and psychiatry. So much has been repeatedly asserted in the literature by some prominent researchers (e.g. Millikan 2002: 114, Garson 2016: ch. 5.3).

Garson (2016) provides a systematic review of the literature on proper functions in biology. The most popular view is the *aetiological* account, also referred to as *selected effects theory*. An early proposal of this sort was offered in Wright (1976: 39, 84ff); many earlier similar accounts are listed in Wright (1976: 94) and Garson (2016: ch. 3.2). Let us spend a few paragraphs discussing this theory and some major objections that have been raised. We will pay more attention to it because, as the reader *in the know* will already have noticed, our account of teleology provided in section III is an aetiological account. On the following paragraphs, we will outline it and defend it against some major criticisms that have levelled at it, without intending to be exhaustive in our defense.

We will follow this discussion by mentioning alternatives to the aetiological account. Due to scope considerations, we fall short of providing an in-depth critical overview of such alternatives. As we have stated, this section's chief intention is to prepare the grounds for Burge's objection to teleofunctionalism, and the fact is that his objection applies to any proposed of teleofunctionalism, be it with the aetiological view of proper functions or with one of the other views we will outline. As such, there is no pressing need to motivate one theory over the others.

The account can be formulated in more than one way, but the general gist is that entities acquire a proper function, that is, a teleology to *do* something, if and only if they have

been *selected* in virtue of doing that something. That is, an entity exists or is arranged in a certain way because it exerted a certain effect which was selected for. We complemented this by stating that the entity selected to do something must have had a *reliable* mechanism to do that something.

Note that we are discussing *natural* proper functions, and therefore the selection process must also be a natural one, as opposed to conscious. Such a selection process would presumably be *causal*, although some recent work on non-causal explanations raises the distant possibility of non-causal and non-mental selection (Lange 2016). To avoid spreading ourselves too thin by exploring all theoretical possibilities, we cofine ourselves to non-mental causal selection processes.

Within aetiological theorists, there is an open debate about what selection processes can institute proper functions. Fine-grained sieves may select small rocks, but this does not imbue small rocks with the teleology to be small! A few other corner examples are discussed in detail in Garson (2016: 51-5). Such discussions appear when we attempt to discern exact necessary and sufficient criteria for our categories. This seldom works out; the world is much too fuzzy. We ourselves believe that there is no hard distinction between selection processes that institute and that do not institute proper functions.

Teleologies, on our account, enter into the picture when non-teleological descriptions fail to capture the counterfactually robust patterns in a complex mechanism—that is, the mechanism's reliability in producing a certain outcome, as with perceptual constancy. Since complexity and robustness are both matters of degree, and in such a way that the eliminability of teleological descriptions is a matter of degree, it follows that bearing a teleology is itself a matter of degree.

Highly teleological systems are ones basically intractable except outside a teleological point of view. That would be the case for modern computers, where the flow of electrons can really only be minimally understood in case we take transistors as *logical gates*, something which depends on the fact that humans design computer hardware to act as logical systems. Other teleological systems may be thinly teleological, so thinly so as not to deserve mention, as exemplified by a rock-sieving process. Our considerations in this paragraph about the attribution of teleology are in the spirit of Dennett (1987)'s considerations relative to the attribution of intentionality.

This gradualist spirit prevents us from getting bogged down on intricate discussions about grey-zone scenarios, either contrived or found in biology. There are no exact criteria for

something constituting evolutionary selection (Godfrey-Smith 2009: ch. 2), and we submit that neither are there exact criteria for selection processes being able to institute teleologies. It is a matter of degree. It is to be determined on a case-by-case bases, in a way respectful of whether scientists have found it necessary to talk of teleology.

To our knowledge, all or almost all examples of natural teleology instituted by selection mechanisms come from causal selection in biology. An most of *those* examples come from evolutionary biology, where causal selection is what is standardly called 'natural selection': differential reproduction based on organismic behavior. These would be called 'biological proper functions' or, for short, *biofunctions*. Biofunctions are the proper functions most associated with teleofunctionalism, so we might call them *teleofunctions* as well.

Proper functions in biology might, in fact, even appear outside the context of evolutionary biology: some state that biofunctions may also be instituted by processes in the very organism itself during its lifetime. One example would be a homeostatic mechanism that could be understood independently of their evolutionary history, being instead defined by a closed feedback mechanism internally selected during the organism's ontogeny (cf. *cybernetic teleology* in Wright 1976: 28). More broadly, any process involving *blind variation* and *selective retention*, as happens with reinforcement learning, would constitute appropriate examples.

A potential example of an ontogenetically-developed proper function would be attributing an antibody with the proper function of detecting and destroying exemplars of SARS-CoV-2.

An individual antibody recruited during a subsequent SARS-CoV-2 infection and which failed to do so could be regarded as unsuccessful. Our understanding of the antibodies during such an infection would be greatly facilitated by casting them in such teleological terms. Other examples of this kind have been offered by Ruth Millikan, David Papineau, and Fred Dretske, among others. Their views are mentioned in Neander (2020: §2) and Garson (2016: ch. 3.4).

In the same section cited above, Garson himself proposes a variation of the aetiological conception of biofunctions which is compatible with ontogenetic acquisition; he calls it the *generalized* selected effects theory. Proper functions can be established not only by differential reproduction—which is the form of causal selection operant in phylogenetic selection and also in some ontogenetic selection—but also by differential reinforcement or persistence, so as to capture more ontogenetic processes. His definition is aimed at processes occurrent in the biological world, either at populations or inside an organism.

The generalized selected effects theory is sufficiently similar to two others theories—the weak aetiological theory (Garson 2016: §6.1) and the systems-based theory (Garson 2016; §6.2)—so that we will not treat the three theories separately. What the weak aetiological theory preconizes is that proper functions attach to inherited traits which contributed to ancestor fitness. It also bears sufficient similarity to the *fitness-contribution theory* which we will review later. The systems-based theory, in turn, states that proper functions attach to traits that ensure their own continued existence during an individual's lifetime, very similar to the ontogenetic aspect of the generalized selected effects theory. Note that the aforementioned systems-based theory is not the famous systems-oriented view proposed in Cummins (1975), which we discuss separately below under the rubric of *causal role theory*.

Having laid out what the aetiological theory of proper functions is, we now consider a few objections based on the summary of criticisms contained in Garson (2016: ch. 3.3). References to the original criticisms can be found in that work. We cannot answer such criticisms in any definite form, for we are not closely acquainted with the standard practices in biology. We describe them briefly and provide sketch answers, so that the reader can evaluate for him or herself. As our intention is not to promote the aetiological theory above other views on proper functions, but instead merely to provoke a critical perspective in the reader, it is not paramount to offer a definite rebuttal of the criticisms below.

Two criticisms have been advanced by Robert Cummins. The first is that (phylogenetically) selected effects cannot explain anything about *individuals*—for example, why any given visual cortex reliably detects color—because individuals are as much a product of their phylogeny as of their ontogeny. Our answer would be that phylogeny robustly predicts and explains individual properties; common ancestry is the reason why members of a species are very similar. We will see this matter again when discussing a third objection below.

The second objection is that selected effects cannot explain whole abilities, such as *seeing*, because at every step along evolutionary history the selection pressure was relative to *small improvements*—at no point there has been selection between seeing and non-seeing animals, but only between one kind of seeing and another kind of seeing. So one couldn't say that wings have been selected for flying. This objection misses the fact that selection pressures can, at every step, point toward a direction; in this example, more accurate vision or better flying. Such a direction is readily identifiable by scientists and figures in successful explanations. This objection again flies in the face of standard scientific practice; it would entail that biologists would be naive to claim that wings were selected for flying. The ease

with which philosophers reject standard scientific practice based on flimsy conceptual arguments can be striking.

A third objection has been raised by Paul Davies. Recall that the posibility of *malfunction* is central to define the notion of proper function. This is because proper function is distinct from what we have called *mere* function, that is, the actual causal profile exhibited by an entity. Something can have a proper function and yet fail to exhibit the associated actual function, thus malfunctioning.

What Davies claims is that the selected effects theory cannot explain normativity. Suppose that a class of objects *A* was selected to do *B*. He argues that the class *A* only contains well-functioning entities, that is, the ones that are in fact able to do *B*. As such, it would be impossible for an *A* entity to fail to do *B*. Therefore, entities selected to perform a function cannot fail to do so. Applied to an example, he would not say that it weren't *hearts* that were selected to pump blood, since the class of all hearts includes disgruntled hearts incapable of blood pumping. If an organism's ancestors had a certain kind of heart, whose good blood-pumping capabilities led such ancestors to pass on their genes, then it was that kind of heart which was selected—and it is incoherent for a heart of the same kind to fail to pump blood.

Our answer is the same as the one provided by Garson (2016: 49): proper function is inherited. Governed by their genome, organisms undergo a systematic developmental procedure which has been selected for to produce hearts, lungs, eyes, and other organs that are observably stable in the species, and which have been selected to perform specific tasks. It is in this systematic context that something can be identified as a *heart*, and it is that structure in the organism that inherits the function to pump blood. Heart *components* may inherit the proper function of executing smaller functions, and malfuctioning hearts may lack those components, but the heart remains a heart and can still be understood as malfunctioning.

We accept the (generalized) selected effects, or *aetiological*, theory of proper functions. However, we have no strong standing on the issue, for four reasons. The first is that, as we have just seen, we do not have a strong grasp in biology so as to be able to defend it from sharp criticism and, thus, promote it over alternative theories of proper function.

The second is that the aetiological theory is a theory of *proper functions*, and there have been researchers who deny that there are any proper functions at all in this teleological sense (see Garson 2016: ch .5). A strong standing on this issue would involve firmly defending the existence of proper functions. This is not something we wish to do, partly

because we have not deeply researched arguments for and against their existence, and mainly because it falls outside the purpose of the present section: to outline *teleofunctionalism*, which *does* presuppose proper functions, so as to better understand Burge's rejection. His rejection is not based on a dismissal of proper functions, but on other considerations entirely. We thus proceed accepting proper functions in biology.

The third and fourth is that the two main conclusions we have derived from it can also be derived in its general form from alternative theories of proper functions, so that choosing between alternatives is not paramount. The first conclusion is that teleologies are systematically related (and in this sense reduced) to non-teleological aspects of the natural order. The second conclusion is that, altough teleological explanations are non-causal, they are constitutively related to their *explanandum*, namely, their own reliable fulfillment. Given that constitutive relation, teleological explanations take the form of organizing principles from which to understand complex mechanical systems.

Absent that constitutive relation, it is hard to see how teleologies can explain anything: either the existence of a trait or the reliable fulfillment of a proper function. We use a theory's ability to make sense of teleological explanations as a criterion of theory choice. However, all available theories of proper function satisfy this *desideratum*.

To show that these two conclusions are congenial to alternative theories of proper function, we will review the two major alternative theories in the order that such alternatives were presented in Garson (2016: chs. 4-5), with only enough detail to show what we intend.

The first theory is the *fitness-contribution* theory of proper function. The proper functions associated to any component of an organism are its *mere* functions which happen to enhance the *present-day* evolutionary fitness of that individual organism or of its encompassing biological population (on average, or some such thing). Fitness-enhancing becomes the purpose of the relevant component of the organism.

The theory can be worked out with a number of distinct details (see Garson 2016: ch. 4.1), and evolutionary fitness is itself a hotly contested notion (Godfrey-Smith 2009: ch. 2), but that is the general gist. This clearly systematically relates proper functions to non-teleological factors, given the available definitions of fitness, which are non-teleological (*idem*), thus adhering to our first conclusion.

What may initially seem to be lacking in this theory is any way to make sense of teleological explanations. How could proper functions attributed *after the fact*—after organisms had already fully developed their fitness-contributing components—explain

anything about the organism? This worry is raised by Garson (2016: 77-8). However, this worry misses the fact that teleologies provide a distinctive sort of explantion: serving as organizing principles. Fitness-based teleologies are constitutively related to their own reliable fulfillment, since such teleologies were instituted only because there *was* a fitness-enhancing mechanism already in place, and can serve as organizing principles to understand how a complicated biological system works. The theory under consideration therefore adheres to our second conclusion as well. As such, we need not choose between the selected effects theory and the fitness-contribution theory of proper functions.

The fitness-contribution theory is sufficiently similar to the *modal theory* reviewed in Garson (2016: ch. 6.3) that we will not consider it separately. The modal theory postulates that a proper function to do so-and-so attaches to traits that contribute to fitness in the closest possible world in which that trait in fact does so-and-so. For example, suppose an animal's wing contains a broken bone or is bereft of feathers, and thus cannot provide lift for flying. The wing will bear the proper function to provide lift for flying just in case flying increases the animal's fitness in the nearest possible world in which that wing *does* provide lift for flying.

The second theory is the *causal role theory* of proper functions. The proper functions associated to any component of an organism are dependant on the proper functions associated to a larger system in which it is a direct sub-component. As such, proper functions are determined from the top-down, leaving open the question of whence comes the original, top-level proper function.

At any rate, the causal role theory of proper functions aims to explain how a typical representative of a top-level system is able of doing complex—something it in fact *does*, regardless of teleology, regardless of what it *should* do—and the theory does it by decomposing the system into sub-systems that play less complex roles, with successive decompositions resulting in simple physical mechanisms. Malfunctional components would deviate from their typical causal contribution to the topmost level causal abilities. This allows for a teleological perspective which helps organize our understanding the system's components components by casting each level of organization, up to somewhat *atomic* components, in terms of a purpose inside an overarching system.

What ends up lacking is a way to understand the purposive behavior of the system as a whole. This has been dealt with differently by different theorists. The analysis provided by Robert Cummins (1975) eschews talk of *proper* functions at the topmost level. The analysis

provided by later theorists, and by further work by Cummins himself, preserve talk of proper functions, but assert them to be subjectively determined by the observer's interest (Garson 2016: 84-5). On neither view do whole systems, such as organisms, bear objective proper funcitons.

Perhaps this is not an issue when discussing perception and mentality in general, since the proper functions associated to these faculties may be understood in terms of their causal contribution to the organism-level capacities understood non-teleologically. If this is so, then the causal role theory can preserve our two conclusions. Teleologies originate from non-teleological causal roles and can serve as ineliminable organizing principles in our theorization about systems otherwise too complex to understand.

To sum up, we have seen how proper functions are central to teleofunctionalist theorization, and how all major theories of proper function are congenial to our proposal about how teleological explantions work.

Burge is in agreement that perceptual psychology is committed to perceptual representations, in a fully teleological sense. Both him and his teleofunctionalist opponents agree that teleological explanation figures centrally in the science. However, he argues that perceptual representational teleology cannot be a biofunction. He claims that these two modalities of teleology are fundamentally incompatible. That is, the success condition which characterizes perceptual representations *cannot*—as a matter of necessity—be to exhibit an ontologically or evolutionarily selected causal profile.

In the next and concluding section, we work out in the best way possible his argument against teleofunctionalism about perceptual representation. We then show what assumptions would allow the teleofunctionalist to resist that conclusion.

Section V. *Against Tyler Burge*

We have reviewed how representational teleologies are systematically related to historical and causal patterns involving the environment and an organism's sensory processing. This has led many researchers, including ourselves, to posit that representational teleologies are constitutively related to such naturalistic patterns. Some among such researchers, again including ourselves, have also argued that teleological explanation would be obscure unless such a constitutive relation obtained.

Burge (2010) emphatically disagrees. What we will call Burge's *central claim* is that the teleologies (success conditions) that characterize a perceptual representation—that is, *representational teleologies*—are necessarily distinct from teleologies instituted by causal selection patterns—that is, *teleofunctions*, which may be *biofunctions* if they were instituted evolutionarily.

The claim is supported in *Origins* by the root mismatch argument (Burge, 2010: 301-3). The present section evaluates that claim and its accompanying argument.

As we have seen, a perceptual representation's success condition is to detect an external condition. The representation must exist if and only if that condition has impacted our senses. We may refer to this fact by saying that representations purport to be *accurate*—they must match an external condition, no matter what. So much is common ground we have already reviewed. Clearly, we concur that representations aim for accuracy.

We have also reached shared conclusions about biological teleofunctions. Whole organisms, or segments thereof, acquire biofunctions through causal selection processes; that is, things are selected due to their causal profile. As a result, biofunctions are related to causal profiles. This is what's called the *selected effects* (or *aetiological*) theory of biofunctions. We will not put this theory in question. However, among aetiological theorists there is room for disagreement on exactly what causal profile that is—on what are the exact contents of a biological teleofunction.

Origins has an identifiable (although not very explicit) stance on this topic. We will claim that the root mismatch argument depends on a particular aetiological theory about the contents of biofunctions, which is implicitly preferred in the book but to which there is a plausible alternative. It is an original contribution of ours to clearly identify what aetiological view the book prefers, how it relates to the root mismatch objection to biological teleofunctionalism, and how it can be avoided through a different aetiological view. The latter is a virtue insofar as it avoids straying too far from Burge's own initial assumptions.

What follows is not explicitly endorsed in the book, although it does seem to be implicitly committed to it. Either way, what is most relevant is that the root mismatch argument depends on it. The view is that, generally, biofunctions pertain to the ultimate goals to which they serve. If a phenotype has been selected by evolution, then its biofunction is to increase fitness, where fitness is a measure of evolutionary success.[6] If it has been selected by ontogenetic processes instead, then its biofunction is to, say increase organismic homeostasis or ensure survival.[7] On either case, biofunctions bear the teleology to aid in *practical* affairs. It is result-driven: what matters is pragmatic success in a given context. Whatever the phenotype does—say, whichever way an organ functions—it has fulfilled its biofunction if and only if an ultimate practical goal has been served, be it fitness increasing or individual survival. We may refer to this fact by stating that biofunctions purport to be *useful* (Burge, 2010: 301).

We can now restate the book's central claim: accuracy and utility are fundamentally distinct. Given the above theory of biofunctions then, should a teleological entity be cast as a biofunction, their success conditions would not be proper to constitute a perceptual representation. Such representations, as is the case for any other teleological entity, are constitutively characterized by their success conditions. The success conditions have to be just right.

Burge offers two arguments for his central, anti-teleofunctionalist claim. We discuss them in turn. Now, one way to construe the central claim is that utility and accuracy are *intensionally* distinct; they have distinct natures (Burge, 2010: 303). Understanding how biofunctions are instituted is insufficient to understand how representations come to be. The teleofunctionalist story is incomplete. By framing Burge's claim in these terms, his first argument can be rendered more clearly.

---

6 *Fitness* is a term of art. There may be no universally applicable concept of fitness (Godfrey-Smith, 2016: ch. 2), but the concepts are not radically distinct. Biologists are skilled at applying the right concept for each case. We may assume that biofunctions acquire the teleology to increase whatever kind of fitness is applicable in a given context.

7 Note that survival is not the same as evolutionary fitness. Survival pertains to individuals; fitness pertains to genomes (roughly put). An individual's sacrifice can further its fitness in case it has aided its genetically-related kins to survive. Conversely, an individual's survival is irrelevant to fitness if it does not lead to its own or its kins' reproduction. This is why fitness is contemporarily referred to as *inclusive* fitness: close relatives are included.

There is a standard way to argue that two entities are intensionally distinct. Namely, to show that it is possible for them to be *extensionally* distinct. They might not be extensionally distinct in the actual world—it might well be the case that obtaining accuracy always leads to increased fitness—but what is relevant is whether it is possible in principle for them to be extensionally distinct. This is precisely Burge's strategy. He aims to show that inaccuracy can increase fitness and that accuracy can decrease fitness. The two are orthogonal to each other. More precisely, it is possible for an organism to fulfill all of its biological functions—every utility-oriented teleology—despite bearing inaccurate representations, proving that accuracy-oriented teleologies are not the same as utility-oriented ones. Equally, he claims, it is possible to fulfill every accuracy-oriented teleology while failing some utility-oriented ones. That is Burge's first argument.

We will present two toy examples. The first example is constructed by a recipe Burge provides for building examples that show how biological and representational functions are intensionally distinct (Burge, 2010: 302).

Let ACTIVE be an organism's bodily state which promotes muscle growth and physical activity; say, it is adrenaline-injected or some such thing. To ensure cardiovascular health and useful muscle toning, but also to avoid excessive energy and nutrient consumption, there is some *optimal frequency* for the organism to exhibit ACTIVE. Sendentarism and hyperactivity are the organism's Scylla and Charybdis. Unlike us, most animals cannot reason their way to discovering what optimal frequency that is. So there must be some naturally selected subconscious mechanism that equilibrates physiological mechanisms which *induce* and which *inhibit* that ACTIVE state. Now suppose that one way in which ACTIVE is induced is when the visual system represents TIGER, prompting a fleeing response. Further suppose that, hitherto, ACTIVE has occurred with sub-optimal frequency. It might just be the case that the best way natural selection finds to boost ACTIVE's frequency is to hijack the visual system so as to increase TIGER's frequency, despite real-world tigers remaining equally frequent. That is, now the visual system was selected to *misdetect* tigers often enough to achieve optimal levels of daily workout. Its biological function is to be accurate only some of the time—since partial accuracy is what best increases fitness. (Burge's particular aetiological view is invoked here.) But the visual system has not stopped representing tigers. The accuracy teleology is still there—but it is not biofunctional.

Here is the second toy example, now coming from speculative theories about the evolution of supernatural thought: the *hypersensitive agency detector*. An agency detector

would aim to detect organisms in general or intentional agents in specific. Its biological utility is clear, ranging from detecting kins and mates to detecting predators and prey. However, as other organisms have evolved to *hide* themselves from predators and prey, agency detection systems are at an arms race with hiding abilities. Quite often, one cannot clearly see an animal but must instead guess its presence from subtle visual and auditory cues. It is hypothesized that this has in fact led many animals to develop a *hypersensitive* agency detector, one which triggers in response to small stimuli which are quite often caused by nothing more than foliage and air currents. Either way, it is surely *possible* for hypersensitive system to evolve. Once again, their biological function would be to be accurate only some of the time; that's how it best promotes fitness. (Again his aetiological view.) But the agency detectors have not thereby stopped representing agents. The accuracy teleology is still there—but it is not biofunctional.

In the two above cases, inaccurate states fulfill all of the perceptual system's biological functions, but still (allegedly) remain as representations. Teleofunctionalist theorists have available to them three responses to Burge's first argument.

First, to claim that not all biological teleofunctions are selected for their causal profile (Vicente, 2012). That constitutes a denial of the aetiological theory of biological functions. We will not explore this response, for it strays too far from Burge's own assumptions.

Second, to deny that the perceptual state in question would really represent something. It has lost its representational capacity, because its biological teleology no longer perfectly matches utility with accuracy. A representational teleology is to detect something (through perceptual constancy), *period*. Not to detect it *some* of the time, but simply to do it. Any misdetection frustrates the representational purpose. We will not explore this response either; it leaves us the teleofunctionalist too vulnerable to the possibility that, as it turns out, we do not perceptually represent many things.

Third, to claim that an organism's part can bear opposing biofunctions at the same time. Above, we implicitly assumed that the organisms originally had a biofunction which matched utility to accuracy—it had a detection biofunction, meant to *always* detect something. The claim here would be that such biofunction would be preserved as the perceptual system was recruited to fulfill other purposes less concerned about accuracy. The new misdetection biofunction—for instance, to induce the ACTIVE bodily state every so often—would exist parallel with the detection biofunction. One can be fulfilled without the other. The detection biofunction remains purely concerned with accuracy. Later, we will show

how this approach can be couched in a systematic aetiological theory of biofunction different from Burge's own. See further discussion in Farias Filho (2018: 82ff).

Enough on the first argument. Burge aims to show that there can be no such thing as a *detection teleofunction*. A central step in his reasoning is the root mismatch claim. We have framed it in intensional terms: utility and accuracy are intensionally distinct. His first argument was to show that they are extensionally distinct, at least in principle.

His second argument is more direct: no mechanism could institute a detection teleofunction. As per his brand of aetiological theory, biofunctions pertain to causal effects on fitness or survival. However, he claims, detection *by itself* is causally irrelevant to these practical purposes. Detection by itself cannot enhance or hinder fitness or survival. Burge's presentation of this point is elliptical (Burge, 2010: 301-2). Below, we present an original framing which renders his point more precise and which allows us to clearly justify his view. However, this gain in precision will also allow us to put forth a significant objection.

What does it mean to claim that detection "by itself" is practically irrelevant? We understand that it means that detection is not *constitutively* associated with any behavioral response (Burge, 2010: 301). Detection is merely correlation instituted through perceptual constancy. Detection requires additional machinery in order to achieve any practical goal and, thus, fulfill any biological teleofuction. As such, detection by itself never fulfills biofunctions. It is only detection *coupled* to a behavioral response that does so.

The intended inference now is that *pure* detection, detection by itself, could never be a biological function of an organism; there can be no pure detection biofunction. Only detection coupled to a behavioral response could be so. However, perceptual representations are characterized by a pure detection teleology. So representations cannot be biofunctions.

The above inference is committed to the principle that all units of natural selection, and thus all bearers of biological teleofunctions, have constitutive, non-coincidental practical effects.

By 'practical effects', we mean effects that are directly selected for or against. Isolated mechanisms such as detection cannot be properly said to be selected for and, thus, to acquire biofunctions. It is only when some behavioral mechanism *consumes* the information in the perceptual system, and thus produces practical effects, that evolutionary or ontogenetic selection can kick in. So it is only the detection-plus-consumption arrangement that can acquire a biological teleofunction. The motivation behind this view is intuitive: evolution and ontogenesis are blind to isolated internal brain events (such as detection mechanisms)

abstracted from their wider consequences on the organism and its behavior. So, properly speaking, they could not ever be selected for.

What we have just stated is the particular aetiological conception of biofunctions to which we have alluded for so many pages now. We might refer to it as *biofunctional holism*, capturing the fact that only composite ensembles are selected for, since only they are constitutively associated with practical effects visible to selection. Isolated organismic components have no biofunctions.

An example from Farias Filho (2018) is illustrative. Suppose chameleons can detect bugs. This ability provides no fitness gains unless it leads chameleons to protrude their tongue and reliably capture the detected bug. That is, detection must be *usefully consumed*. Biofunctions are attached to the detection-plus-capture ensemble, not to detection itself. Biological success is achieved if and only if fitness is increased. Should a chameleon *misdetect* a cherry as a bug, it will still capture nutritious food and fulfill its biofunctions. Should a chameleon *detect* a bug but fail to capture it, it will thoroughly fail its biological purposes. This demonstrates the accuracy-utility mismatch.

Biofunctional holism has broad implications. Applied thoroughly, the detection-plus-capture system in chameleons is not yet the locus of biofunctions. There is no fitness gain from capturing a bug unless the chameleon's digestive system is in order, so as to extract energy from the bug. Furthermore, such energy is irrelevant unless the chameleon's motor systems are in order, so as to use that energy for fleeing and mating. In the end, its reproductive systems must be in order too; the chameleon can eat, flee, and mate all it wants, its fitness will remain *zero* if it cannot have offspring.

The upshot is that it is only *the whole chameleon* (or something close) which is associated with a biofunction. Its biofunctions are fulfilled if and only if the *concerted activity of the whole organism* increases fitness. This conclusion is very strong and, yet, we see no principle weaker than biofunctional holism which can justify Burge's claim. Recall our rendition of his claim: since detection is *not* constitutively linked to practical effects, detection cannot be a biofunction. However, as we have shown, only whole organisms are constitutively linked to practical effects, such as survival or fitness. Without noticing, Burge committed to a very strong consequence.

We believe he would not accept it. Thankfully, the consequence can be avoided through an alternative aetiological theory of perception. We may call the alternative *biofunctional modularism*. We must assume that organisms have at least a quasi-modular

organization, so that they contain some clearly discernible subsystems σ_1 through σ_N. Consider a successfully selected organism. Suppose that during selection each of its subsystems σ_K had a small causal role, some localized mechanical function F_K. Since the organism was selected, what happened is that the concerted activity $\langle F_1, \ldots, F_N \rangle$ of its modules increased fitness during selection.

Biofunctional holism would have that the whole organism—the ensemble $\langle \sigma_1, \ldots, \sigma_N \rangle$—acquired the biological function to increase fitness, regardless of how that is done. The ensemble would fulfill its function even if it increased fitness through an entirely distinct functional network $\langle F^*_1, \ldots, F^*_N \rangle$.

Biofunctional modularism would have it that each σ_K acquired the biological function to do F_K—that which it was doing during selection, that small isolated contribution. At any given occasion, it is irrelevant whether doing F_K would really increase fitness. Its biofunction would be fulfilled just in case it did F_K. Each quasi-modular subsystem in our chameleon's body would acquire distinctive biofunctions pertaining to what it was doing when the chameleon evolved and regulated itself. Independently of anything else, a chameleon's tongue has the biofunction to protrude, its stomach the biofunction to digest and, most crucially, *its visual system the biofunction to detect bugs*, since bug detection was indeed a factor in the chameleon's success. Biofunctional modularism shows a clear way in which detection biofunctions can originate.

Furthermore, it allows for plausible responses to both of Burge's arguments. The second argument was directly answered above: it is not just because detection *by itself* has no practical effects that it cannot partake in a functional network which does and, thereby, acquire biofunctions related to its own isolated detection task. As to the first argument, there seems to be no problem in accepting that the same organismic module can acquire two biofunctions, since that module can at the same time partake in different functional networks each of which increase fitness. As a result, it can well acquire the biofunction to detect external objects—bugs, tigers, agents—since that is clearly useful on so many occasions, as well as the biofunction to misfire some of the time so as to serve other purposes. All that this entails is that the module might be unable to fulfill all of its biofunctions at once. But detection biofunctions can sure co-exist with hypersensitive agent detectors and the like.

Biofunctional holism has broad implications which Burge would certainly wish to avoid. Unless the dichotomy between modularism and holism is false, this constitutes a strong reason to embrace biofunctional modularism. This, however, invalidates the two arguments

for the root mismatch claim which we reconstructed from Burge's work. As such, biological teleofunctionalism about perceptual representation survives in a way compatible with perceptual psychology.

Conclusion

Following the scientific reconstruction in *Origins of Objectivity* (2010), on Section I we have studied the basic representations of perception. These are perceptual representations as of type attributes. The teleological component of representational explanation is seen as ineliminable in the science, according to which perceptual representations are subject to evaluative norms pertaining to what type attribute has caused them and through what internal processing mechanism. Teleology is what marks the distinction between merely mechanical processing and genuine perception, as further explored in Section III.

Section II reviewed our empirical knowledge of how perceptual constancy works, highlighting how many points of dissent co-exist with a broad consensus on its general aspects, defending Burge's view against empirical criticism. We have seen how contemporary accounts of perceptual constancy may solve two age-old philosophical-cum-scientific conumdrums, the distality problem and the underdetermination problem, whose solutions are paramount to Burge's philosophy of perception.

Having checked that Burge's rendition of perceptual constancy is accurate, Section III mobilized that notion to explain *why* perceptual psychology became irretrievably committed to representations: perceptual constancy makes it so perception cannot be understood systematically through inner signal processing, but only by representations and representation transformations. Should we be confident that the science is accurate, this provides strong empirical reason to accept some form of representationalism and all its philosophical consequences (whatever they are).

Section IV argued that biological systems acquire a representational teleology if and only if they have been *mechanically selected* to detect such an attribute through a perceptual constancy mechanism. That is, we argued for biological teleofunctionalism about perceptual representation.

In Section V we had to fend off Tyler Burge's reasons to believe that no selection process could institute representations. His reasons were implicitly committed to holism about biofunctions,

a thesis with unpalatable results; we thus created an alternative modularist approach to biofunctions which avoids such results but also invalidates Burge's reasoning, thus protecting teleofunctionalism from his notorious objection and thus rendering is *prima facie* compatible with perceptual psychology.

Although empirical science cannot substitute philosophical work, history has shown how much of a guide it can be. We hope to have contributed in this work to humanity's understanding of the foundations of mind, a source of long-lasting puzzlement for intellectuals.

Bibliography

— ALLEN, C. & NEAL, J. “Teleological Notions in Biology.” In: ZALTA, Edward (ed.). *Stanford Encylopedia of Philosophy*, 2020.

— BEGBY, Endre. “Origins of Objectivity.” *Notre Dame Philosophical Reviews*, 2011.

— BLOCK, Ned. “Seeing-As in the Light of Vision Science.” *Philosophy and Phenomenological Research*, v. 89, n. 3, pp. 560-72, 2014.

— BONJOUR, Laurence. “In Search of Direct Realism.” *Philosophy and Phenomenological Research*, v. 69, n. 2, pp. 349-67, 2004.

— BRAITENBERG, Valentino. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press, 1986.

— BRISCOE, R. & GRUSH, R. “Action-based Theories of Perception.” In: ZALTA, Edward (ed.). *Stanford Encyclopedia of Philosophy*, 2015.

— BROWN, Curtis. “Narrow Mental Content.” In: ZALTA, Edward (ed.). *Stanford Encyclopedia of Philosophy*, 2016.

— BROWN, Richard. “ Background and Illuminants: The Yin and Yang of Colour Constancy.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003, ch. 8.

— BURGE, Tyler. “Perceptual Entitlement.” *Philosophy and Phenomenological Research*, v. 67,

n. 3, pp. 503-48, 2003a.

— BURGE, Tyler. “Some Reflections on Scepticism: Reply to Stroud.” In: HAHN, M. & RAMBERG, B. (eds.). *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, MA: MIT Press, 2003.

— BURGE, Tyler. “Disjunctivism and Perceptual Psychology.” *Philosophical Topics*, v. 33, n. 1, 2005.

— BURGE, Tyler. “Primitive Agency and Natural Norms.” *Philosophy and Phenomenological Research*, v. 79, n. 2, 2009a.

— BURGE, Tyler. “Five Theses on De Re States and Attitudes.” In: ALMOG, J. & LEONARDI, P. (eds.). *The Philosophy of David Kaplan*. Oxford, UK: Oxford University Press, 2009b, pp. 246-324.

— BURGE, Tyler. *Origins of Objectivity*. Oxford, UK: Oxford University Press, 2010.

— BURGE, Tyler. “Perception: Where Mind Begins.” *Philosophy*, v. 89. n. 3, 2014.

— CAMPBELL, John. “Origins of Objectivity by Tyler Burge.” *The Journal of Philosophy*, v. 108, n. 5, pp. 269-85, 2011.

— DENNETT, Daniel. *The Intentional Stance*. Cambridge, MA: MIT Press, 1987.

— FARIAS, Sérgio. *Naturalising Intentionality: A Teleological Approach*. PhD Thesis in Philosophy, King’s College London, 2018.

— FRISBY, J. & STONE, J. *Seeing: The Computational Approach to Biological Vision*. Cambridge, MA: MIT Press, 2010.

— FODOR, Jerry. “A Theory of Content, I: The Problem.” In: FODOR, Jerry. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press, 1990.

— FODOR, Jerry. “Burge on Perception.” In: MARGOLIS, E.; LAURENCE, S. (eds.). *The Conceptual Mind:* New Direction in the Study of Concepts. Cambridge, MA: MIT Press, 2015, ch. 8.

— GARSON, Justin. *A Critical Overview of Biological Functions*. New York, NY: Springer, 2016.

— GILCHRIST, Alan. “The Importance of Errors in Perception.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003, ch. 14.

— GODFREY-SMITH, Peter. *Darwinian Populations and Natural Selection*. Oxford, UK: Oxford University Press, 2009.

— HOFFMAN, Donald. “The Interaction of Colour and Motion.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003a, ch. 12.

— HOFFMAN, Donald. “An Internalist Account of Color.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003b, commentary on ch. 13.

— HILBERT, David. “Constancy, Content, and Inference.” In: HATFIELD, G. & ALLRED, S. (eds.). *Visual Experience: Sensation, Cognition, and Constancy*. Oxford, UK: Oxford University Press, 2012, ch. 10.

— LANGE, Marc. *Because Without Cause: Non-Causal Explanations in Science and Mathematics*. Oxford, UK: Oxford University Press, 2016.

— LUNTLEY, Michael. *Contemporary Philosophy of Thought:* Truth, World, Content. Oxford, UK: Blackwell, 1999.

— MACLEOD, D. & GOLZ, J. “A Computational Analysis of Colour Constancy”. In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003, ch. 7.

— MILLIKAN, Ruth. “Biofunctions: Two Paradigms.” In: ARIEW, A. *et al.* (eds.). *Functions: New Essays in the Philosophy of Psychology and Biology*.

— NEADER, K. & SCHULTE, P. “Teleological Theories of Mental Content.” In: ZALDA, Edward (ed.). *Stanford Encyclopedia of Philosophy*, 2020.

— PALMER, Stephen. *Vision Science:* Photons to Phenomenology. Cambridge, MA: MIT Press, 1999.

— RESCORLA, Michael. “Perceptual Constancies and Perceptual Modes of Presentation.” *Philosophy and Phenomenological Research*, vol. 88, n. 2, pp. 468-76, 2014.

— SCHULTE, Peter. “Perceiving the World Outside: How to Solve the Distality Problem for Informational Teleosemantics.” *Philosophical Quarterly*, v. 68, n. 271, 2018.

— SCHULTE, Peter. “The Nature of Perceptual Constancies.” *Philosophy and Phenomenological Research*, v. 103, n. 1, 2021.

— SHEBILSKE, W. & PETERS, A. “Perceptual Contancies: Analysis and Synthesis”. In: PRINZ, W. & BRIDGEMAN, B. (eds.). *Handbook of Perception and Action, Volume 1*, 1996, ch. 7.

— STONE, James. *Vision and Brain:* How We Perceive the World. Cambridge, MA: MIT Press, 2011.

— WAGNER, Mark. “Sensory and Cognitive Explanations for a Century of Size Constancy Research.” In: HATFIELD, G. & ALLRED, S. (eds.). *Visual Experience: Sensation, Cognition, and Constancy*. Oxford, UK: Oxford University Press, 2012, ch. 3.

— WEBSTER, Michael. “Light Adaptation, Contrast Adaptation, and Human Colour Vision.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003, ch. 2.

— WHITTLE, Paul. “Contrast Colours.” In: MAUSFELD, R. & HEYER, D. (eds.). *Colour Perception: Mind and the Physical World*. Oxford, UK: Oxford University Press, 2003, ch. 3.

— WILCOX, Stephen. “Functionalism: Then and Now.” In: OWENS, D. & WAGNER, M. (eds.). *Progress in Modern Psychology: The Legacy of American Functionalism*, 1992, ch. 3.

– WRIGHT, Larry. *Teleological Explanations: An Etiological Analysis of Goals and Functions*. Berkeley, CA: University of California Press, 1976.