

CONSORZIO
di DOTTORATO
in FILOSOFIA
NORD-OVEST


An Analysis of the Concept of Consciousness

LUIS GUILLERMO CASTRO MARTIN

luisguillermo.castromartin@unito.it

A Dissertation Submitted to the
Northwestern Italian Philosophy Consortium (FINO)
In Partial Fulfillment of the Requirements for the Degree of
DOCTOR OF PHILOSOPHY
In the Department of Philosophy
UNIVERSITÀ DEGLI STUDI DI GENOVA

2021



Title of Dissertation/Titolo della Dissertazione:

An Analysis of the Concept of Consciousness (Un'analisi del concetto di coscienza)

Author/Autore:

Luis Guillermo Castro Martin

Year/Anno:

2021

Examination Committee/Commissione di esame:

- **Alberto Voltolini (Thesis Supervisor, University of Turin)**
- **Sam Coleman (University of Hertfordshire)**
- **Luke Roelofs (New York University)**
- **Alfredo Paternoster (University of Bergamo)**

The copyright of this Dissertation rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

End User Agreement

This work is licensed under a Creative Commons Attribution-Non-Commercial-No-Derivatives 4.0

International License: <https://creativecommons.org/licenses/by-nd/4.0/legalcode>

You are free to share, to copy, distribute and transmit the work under the following conditions:

- *Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).*
- *Non-Commercial: You may not use this work for commercial purposes.*
- *No Derivative Works: You may not alter, transform, or build upon this work, without proper citation and acknowledgement of the source.*



In case the dissertation would have found to infringe the policy of plagiarism it will be immediately expunged from the site of FINO Doctoral Consortium.



UNIVERSITÀ
DEGLI STUDI
DI TORINO



UNIVERSITÀ
DEGLI STUDI
DI PAVIA

Abstract

The aim of this research is to introduce an interpretation of the concept of consciousness that is metaphysically compatible with most contemporary theories of consciousness, insofar as it is intended to fix the conceptual level to which the study of consciousness belongs. In the first chapter, I analyze different uses of the terms 'awareness' and 'consciousness' and three approaches to the concept of consciousness that can be drawn from three popular kinds of theories of consciousness. My goal is to show that those approaches are based on the identification of three discernible features of conscious states, namely, intentionality, reflexivity and phenomenality. In the second chapter, I introduce an interpretation of the concept of consciousness where the above features are characterized as components of the concept, unifying them into a single conceptual space. Such interpretation allows me to explore how the conceptual space can be adjusted to build models of consciousness. In the third chapter, I argue that conscious states are internal states that are described in terms of the components of the concept of consciousness and their relations. In the fourth chapter, I show how my view can be implemented to theorize about non-human consciousness by analyzing two kinds of non-human systems, namely, dolphins and bees.

Keywords: conceptual space, practice of conceptualization, discernible feature, intentionality, reflexivity, phenomenality.

Acknowledgements

There are many people to thank. I would like to start by thanking my supervisor Alberto Voltolini for his support since the beginning of this journey. I would not have been able to put my ideas into words without his commitment and professionalism. I would also like to thank Luke Roelofs, Sam Coleman, Alfredo Paternoster, Carlo Penco, Andrea Iacona, Tim Bayne, Wanja Wiese, Luca Pomilio, Jesús Baceta, Leopoldo Márquez and Miguel Vásquez for their comments and guidance. I also thank my brother, Carlos Castro, for helping me with the language and style of this work. I would like to specially thank my wife, Glorymar Hernández, for her patience and motivation; I owe much of this research to the fruitful discussions we had during the last three years. Finally, I thank my family and friends for their encouragement.

For Glorymar

Table of Figures

Figure 1: The Conceptual Space of Consciousness	33
Figure 2: Ways to Describe Internal States	45
Figure 3: Intentional States	53
Figure 4: Reflexive States	59
Figure 5: Phenomenal States	65

Table of Contents

Introduction	1
Chapter 1: Ways to Define Consciousness.....	5
1.1. Ordinary uses of ‘awareness’ and ‘consciousness’	7
1.2. Technical uses of ‘awareness’ and ‘consciousness’	10
1.3. Three approaches to consciousness	13
1.3.1. Intentionality.....	15
1.3.2. Reflexivity	19
1.3.3. Phenomenality	24
Chapter 2: The Concept of Consciousness.....	29
2.1. An interpretation of the concept of consciousness	30
2.2. Three models of consciousness	39
Chapter 3: Conscious States.....	44
3.1. Intentional states.....	45
3.2. Reflexive states.....	54
3.3. Phenomenal states	59
Chapter 4: Characterizing Non-Human Consciousness.....	67
4.1. Dolphin consciousness	70
4.1.1. The dolphin world	71
4.1.2. The social dolphin	77
4.1.3. Can we attribute consciousness to dolphins?	80
4.2. Bee consciousness	85
4.2.1. Beyond traditional ethology	87
4.2.2. Bee emotions?.....	94
4.2.3. Can we attribute consciousness to bees?	97
Concluding Remarks	99
Bibliography	102

Introduction

To explain consciousness is one of the biggest scientific challenges of our time, yet it seems to be something with which we are very well acquainted: we have felt headaches, entertained thoughts, seen objects, heard sounds, etc. In a broad sense, we could all agree that consciousness is experience, i.e., a particular mode of interacting with ourselves and the external world that contrasts with purely mechanical things, like watches and windmills. However, there does not seem to be a clear way to explain why experience is the way it is and how it emerges. Moreover, although we consider ourselves the paradigm of conscious beings, we frequently attribute at least some aspects of consciousness to non-human beings, like dogs and elephants, which means that consciousness might not be an exclusively human trait.

After centuries of scientific reasoning and experimentation, there is no definite explanation of how and why consciousness emerges from, say, a bunch of interconnected neurons. No matter how close we look, we cannot seem to find the sweetness of an apple or the mental representation of a cat in our brains. From the theoretical point of view, explaining how consciousness emerges involves answering a set of questions called the *easy problems*, while explaining why it emerges is considered the *hard problem of consciousness* (Chalmers, 1995). Although it is reasonable to assume that, at some point, we will have the tools to explain the easy problems, theorists often point out that the hard problem will still remain, insofar as an explanation of the neural mechanisms that underly consciousness is not an explanation of why experiences are the way they are. Thus, we find ourselves in a

situation where there are many plausible theories of consciousness and no clear reason to rule out any of them.

A particularly challenging fact about consciousness is that, in contrast to most scientific subject-matters, it cannot be directly measured nor observed; consider that I am only acquainted with my own experiences and that I can only infer what you are experiencing under certain circumstances. In fact, I can only affirm that you are a conscious being indirectly, for instance, by means of interpreting your behavior and brain activity as signs of your having conscious states. For this reason, many theorists have argued that the study of consciousness requires a new methodological paradigm. Although I concur with this observation, I think that the reason why we are in need of a different methodology is that the study of consciousness belongs to an irreducible conceptual level.

Since I cannot decide where I stand regarding the debate on the nature of consciousness, I will attempt an alternative approach. Instead of providing reasons to endorse a kind of theory over another, I will analyze our understanding and use of the concept of consciousness, hoping to provide a way of studying it that avoids some of the most pressing issues concerning its nature. Hence, the interpretation of the concept of consciousness that I will present here does not directly oppose any particular theory of consciousness because it says nothing about its nature; this is the reason why I will hold that my proposal is metaphysically compatible with most current theories of consciousness.

My interpretation of the concept of consciousness can be seen as a sort of pragmatic approach to the study of conscious systems, similar to Dennett's (1971) stance-based account of intentional systems. To paraphrase Dennett's words (p. 87), I will hold that "a particular thing is a [conscious] system only in relation to the strategies of someone who is trying to explain and predict its behavior". In consequence, my interpretation of the concept of consciousness will not imply that current theories of consciousness are mistaken, but that adopting a conceptual approach enables the identification of what all of them have in common, namely, they are different ways to interpret a single conceptual space. Hence, my approach is not only metaphysically compatible with most contemporary theories of consciousness, but also with different views on the ultimately correct analysis of the concept of consciousness¹, insofar as it admits the customization of the conceptual space according to our theoretical interests or strategies.

¹ I thank Luke Roelofs for calling my attention to this characteristic of my proposal.

In support of this conceptual approach, consider that the history of science is full of cases where the proper characterization of a concept has led to a satisfactory explanation of a certain phenomenon. Analogously, I will suggest that a different characterization of the concept of consciousness might lead to new ways of theorizing about it. As Sheldrake points out, “[m]any scientific concepts –from time to chemical bonds to genes to species– lack stable definitions but remain helpful categories to think with” (2020, pp. 17–18). Thus, although the term ‘consciousness’ seems to lack a stable definition, my plan is to build a case for Chalmers’s suggestion that “the concept of consciousness is irreducible, being characterizable only in terms of concepts that themselves involve consciousness” (1996, p. 106) In particular, I will argue that a proper characterization of the concept of consciousness has to be given in terms of the concepts of intentionality, reflexivity and phenomenality.

In the first chapter, I analyze different uses of the terms ‘awareness’ and ‘consciousness’, which enables the conceptual distinction between the functional basis of consciousness and consciousness itself, as well as three ways to define consciousness that are based on the identification and conceptualization of three paradigmatic features exhibited by conscious states, namely, intentionality, reflexivity and phenomenality. My suggestion is that those ways to define consciousness correspond to different theoretical strategies or, as I rather call them, practices of conceptualization. Accordingly, I argue that, since intentionality, reflexivity and phenomenality are discernible features of conscious states, or aspects of consciousness, their concepts point to three characters exhibited by some internal states.

In the second chapter, I introduce an interpretation of the concept of consciousness that combines the concepts of intentionality, reflexivity and phenomenality to form a three-dimensional conceptual space. The purpose of this interpretation is to highlight the conceptual interdependence between the concept of consciousness and the concepts of its features. Then, I explore three general ways to reconfigure the conceptual space to build models of consciousness. My goal is to show that there are different ways to interpret said conceptual space that depend on our theoretical interests, not different concepts of consciousness.

In the third chapter, I rely on my interpretation of the concept of consciousness to present a way to characterize conscious states. I start by elaborating on my claim that the concepts of intentionality, reflexivity and phenomenality point to three paradigmatic characters exhibited by some internal states to claim that conscious states are descriptions of

those internal states. Thus, I introduce a classification of conscious states in terms of three general categories, namely, intentional, reflexive and phenomenal states. The aim of this chapter is to argue that our attributions of consciousness rely on our ability to identify and conceptualize the characters exhibited by some internal states.

In the fourth chapter, I employ my interpretation of the concept of consciousness to assess two possible kinds of conscious systems, namely, dolphins and bees. My goal is to illustrate how my view can be implemented to theorize about non-human consciousness. The methodology that I adopt there is that of interpreting recent scientific observations of dolphin and bee behavior in terms of my interpretation of the concept of consciousness. My aim is to show that the concept of consciousness can be modelled according to the specific characters exhibited by the internal states of each kind of system, thus producing a conceptualization of the kind of consciousness that the system appears to entertain. Hence, what I call *dolphin consciousness* and *bee consciousness* are different configurations of the conceptual space of the concept of consciousness, not independent concepts of consciousness.

It should be clear from the outset that this research is about a way to approach consciousness that does not depend on taking a stand regarding its ontological status. Consequently, my view does not oppose any current theory of consciousness because it says nothing about its ultimate nature. My intention is to present a conceptual background that merely relies on the conceptual space of the concept of consciousness to support the hypothesis that the study of consciousness belongs to an irreducible conceptual level. Hence, regardless of the ontological status of consciousness or the ultimate characterization of its concept, my suggestion is that consciousness must be studied in its own terms.

Chapter 1: Ways to Define Consciousness

*Defining consciousness is problematical
not just because of the elusive, will-o-the-wisp nature of the
phenomenon itself but because we can mean so many different things
by the one word.*

Marian Stamp Dawkins (1998, p. 4)

To most of us, it is clear that consciousness is a real phenomenon. However, there is no general agreement on how to characterize its nature or its concept. In scientific contexts, the term ‘consciousness’ is often a synonym of ‘awareness’ that denotes a function or set of functions carried out by the brain; in philosophical contexts, the term may refer, among many other alternatives, to a substance, a property, a feature, a character, or a mode. Moreover, the study of consciousness is often associated to a series of perplexing theoretical problems within Western philosophy and science², such as the mind-body problem, the problem of other minds, the binding problem, the palette problem, and the explanatory gap.

In the last few decades, theorists have proposed several ways to explain the nature of consciousness. Some have argued that it is part of the natural world (Chalmers, 1996), others have tried to explain it in terms of something else (Baars, 1988), and still others have claimed that there is no place for it in our scientific endeavors (Churchland, 1995). All of these

² This does not mean that explaining consciousness is exclusively a Western issue. What it means is that the contemporary study of consciousness originated from the Western interpretation of human nature; that consciousness is nowadays investigated worldwide is a consequence of the globalization of academic education. In contrast, there are cultures that characterize consciousness in ways that prevent its study from the perspective of standard academic science and philosophy (Maffie, 2008; Tinker, 2004).

metaphysical alternatives somehow determine the nature of consciousness, insofar as they fix its ontological status. From the philosophical perspective, most theorists either adopt a dualistic or a monistic ontology, thus creating an independent realm for consciousness, placing it inside the natural world, or eliminating it from reality. According to a further group of theorists, like McGinn (1991b) and Pinker (2009), the ultimate nature of consciousness is beyond our reach.

In this chapter, I will adopt a different approach, namely, I will examine how we use the term ‘consciousness’ to examine our understanding of the concept of consciousness. Thus, the starting point of this chapter, and the main premise of this whole research, somehow opposes the usual way to approach consciousness: I will assume that, since the term ‘consciousness’ can mean many different things, perhaps we are not really acquainted with its referent, i.e., maybe we do not know what consciousness is.

To avoid misinterpretations, I will refer to conscious things as systems, instead of using the more frequent terms ‘subject’ or ‘mind’. The term ‘system’ does not presuppose the object/subject distinction and it does not imply that conscious things have to be organisms or living things. Conceptually speaking, we should not exclude the possibility of there being conscious systems that do not share our neurobiology or the behavioral responses typically associated to consciousness; contemporary ethological research (Allen & Bekoff, 1999; Stamp Dawkins, 1998) suggests that this is not a mere possibility, but a fact. As a methodological tool, I will use simple quotation marks for words and terms, italics³ for concepts, and regular text for their referents. Thus,

- ‘consciousness’ ⇒ (the English word)
- *consciousness* ⇒ the concept of consciousness
- consciousness ⇒ real/objective consciousness, i.e., the referent of *consciousness*.

In a nutshell, the aim of this chapter is to explicate, in Carnap’s sense (1950, pp. 3–8), our standard uses of ‘consciousness’. In the first two sections, I will analyze both ordinary and technical uses of ‘awareness’ and ‘consciousness’ to show that *awareness* is not identical to *consciousness*. In the last section, I will discuss three possible ways to characterize *consciousness* that can be drawn from contemporary theories to argue that theorists tend to define the concept by identifying three discernible features of conscious states, namely, intentionality, reflexivity and phenomenality.

³ As usual, I will also use italics to make emphasis and to name views and theories.

1.1. Ordinary uses of ‘awareness’ and ‘consciousness’

Since human cultures have a longstanding propensity to anthropomorphize certain features of the world (Daston & Mitman, 2007), maybe we should start by distinguishing standard uses of ‘consciousness’ and ‘awareness’ from their metaphorical uses, as when we say “such-and-such company is eco-conscious”. In this context, some theorists have assessed whether “ordinary people are disposed to ascribe different mental states to entities that are given behaviorally and functionally equivalent descriptions” (Sytsma & Machery, 2009, p. 21), and showed that our ordinary use of mental terms is considerably vague. Similarly, others have proposed that “culturally characteristic environments may afford distinctive patterns of perception” (Miyamoto et al., 2006, p. 113), which suggests that our attribution of certain properties to things in the world, and even to ourselves, might be influenced by how our different social environments affect our perceptual and cognitive processes. Throughout this chapter, I will argue in favor of the hypothesis that our attributions of awareness and consciousness are determined by what I call *practices of conceptualization*⁴, which are conventions that concern our understanding of *consciousness* and *awareness*.

The terms ‘consciousness’ and ‘awareness’, as well as their cognates, are frequently used by native speakers of several Western languages⁵, so someone might object that their ordinary meanings are relatively clear, and that the difficulty lies in providing scientifically useful definitions of them. However, consider that, although both terms are often taken to be synonyms, they are not always interchangeable, or at least not in English. In everyday language we find expressions like “such-and-such company is eco-conscious” or “my cat was not aware of the glass door when it ran into it”, but it would be odd to say, “my dog was consciously eating my homework” or “my newborn cousin is aware of global warming”. I take this asymmetry of use as a first sign that *consciousness* and *awareness* are not the same concept.

⁴ By ‘practice of conceptualization’ I intend to accentuate the fact that the attribution of consciousness, or awareness, is something that we put into practice once we have learned to use ‘consciousness’, or ‘awareness’, and that such practice is conditioned by cultural and theoretical parameters that are constantly changing in accordance with new scientific discoveries and historical events. Put briefly, practices of conceptualization are explanatory strategies that we adopt according to specifiable theoretical goals. Since these practices are shaped by environmental, perceptual and cognitive factors that are beyond the conceptual scope of my work, I will not say anything about how they are acquired.

⁵ I will focus on uses of these terms in English, but the situation is similar in other Western languages. For instance, although the Italian word ‘consapevolezza’ is frequently translated as ‘awareness’ and ‘coscienza’ as ‘consciousness’, they are often interchangeable. In Spanish, while ‘awareness’ and ‘consciousness’ are both translated as ‘consciencia’, their meanings might vary according to the context.

Let us now review how we ordinarily use the terms ‘awareness’ and ‘consciousness’. The online version of the *Oxford English Dictionary* (OED) defines the noun ‘awareness’ as “[t]he quality or state of being aware, consciousness; (also) the condition of being aware (of something or that something is)”. Likewise, the adjective ‘aware’ refers to something that is “[i]nformed, cognizant, conscious, sensible”. When people say that a system is aware of something or that something is the case, they mean that the system is actually perceiving, thinking, imagining, etc., a certain object or state of affairs. Thus, our everyday use of ‘awareness’ is associated with being awake, having knowledge of something, or paying attention to something.

Broadly speaking, we can distinguish an intransitive and a transitive use of ‘awareness’: when I wake up in the morning, I enter a general state of awareness; when I hear certain chimes, I become aware that my phone is ringing. In its intransitive use, “being aware” is equivalent to “being in a certain psychological state”. In its transitive use, noticing the glass door at the entrance of the office is often described as becoming aware of it. It is noteworthy that, while intransitive awareness conflates with intransitive consciousness, as in “the patient recovered awareness/consciousness”, transitive awareness rarely conflates with transitive consciousness: being aware of the negative effects of cocaine is not exactly the same as being conscious of them. Finally, in everyday language, the transitive use of ‘awareness’ is more common than the intransitive one; consider that, if I enter the room and say, “I’m aware”, you will probably ask, “what is it you’re aware of?”

Therefore, when people say that a system has awareness, they usually refer to the system’s ability to attend to its environment and to itself, i.e., in the transitive sense of the term. In contrast, the intransitive use of ‘awareness’ points to the general state of being awake, though not necessarily attentive. Imagine that, while reading this sentence, someone comes in the room and asks, “are you aware that the window is open?” Notice that the question is not about whether you can see the window open, but about whether you know that it is. Similarly, when you wonder whether you closed the refrigerator door before leaving the kitchen, it is because you are not aware of having done it; when you turn back and see that it is closed, or remember doing it, you become aware of having closed it. Notice that, in these last two examples, a native speaker of English would probably abstain from replacing the term ‘aware’ with ‘conscious’: being conscious about leaving the refrigerator open points to a further fact, perhaps that you are wasting electricity.

Let us now analyze the ordinary use of the term ‘consciousness’ and its cognates. The online version of the OED gives the noun ‘consciousness’ six different definitions, the first of which says: “internal knowledge or conviction; the state or fact of being mentally conscious or aware of something”. Notice that ‘conscious’ and ‘aware’ are here taken as synonyms. The rest of the OED’s definitions of ‘consciousness’ either point to derivative uses of this first definition or to uses that apply to different contexts. As for the adjective ‘conscious’, the OED gives eleven different definitions. The first says that to be conscious amounts to “having awareness of one’s own wrongdoing, affected by a feeling of guilt”, while the second one reads, “having knowledge or awareness; able to perceive or experience something”. Notice that the first definition applies to moral contexts, while the second points to the ability to experience, which is the sense that interests us.

Broadly speaking, most of our everyday uses of the term ‘consciousness’, and its cognates, refer to moral awareness. Maybe I am aware of having left the door of the refrigerator open without considering its environmental consequences, in which case I would not be morally conscious of my action. People who drive under the influence of alcohol or drugs might be aware of the red light, but probably they are not conscious of how dangerous it is to pass through it. However, it is the less common use of the term ‘consciousness’, the one that refers to experience, what motivates the philosophical debate. According to many theorists, some states of awareness are accompanied by a certain what-it’s-like-ness, though not all of them: sometimes soldiers are aware of their wounds without being conscious of the pain they cause them, and sometimes they feel an excruciating pain without being aware of its cause. This sense of the term ‘consciousness’ is the second sign that *consciousness* is not identical to *awareness*.

This brief analysis has shown that our everyday use of ‘consciousness’ cannot be detached from that of ‘awareness’; in fact, the OED defines ‘consciousness’ in terms of awareness. Hence, it seems that *consciousness* and *awareness* must be somehow intertwined. However, we have also seen that the terms are not always interchangeable, which means that ‘consciousness’ and ‘awareness’ do not point to the same concept. In the following section, I will try to identify the difference between *consciousness* and *awareness* by analyzing how theorists tend to use the terms ‘awareness’ and ‘consciousness’.

1.2. Technical uses of ‘awareness’ and ‘consciousness’

There is no generally agreed definition of consciousness. As Van Gulick (2018) says, “[t]he words ‘conscious’ and ‘consciousness’ are umbrella terms that cover a wide variety of mental phenomena”. Although in everyday language ‘awareness’ and ‘consciousness’ coincide in many of their uses, it seems that ‘awareness’ points to a certain cognitive ability of some systems, while ‘consciousness’ points to the ability to experience. For instance, Chalmers says that “[c]onsciousness is always accompanied by awareness, but awareness [...] need not be accompanied by consciousness” because “[o]ne can be aware of a fact without any particular associated phenomenal experience” (1996, p. 28). Recall that this research is not about consciousness, so I will not take a stand regarding this claim. In this section, I will explore the conceptual reasons that might justify this kind of view, which will serve to identify the theoretically relevant sense of the term ‘consciousness’.

While the use of the term ‘awareness’ in ordinary contexts seems fairly clear, this is not the case in philosophical and scientific contexts. Despite the fact that *awareness* plays a fundamental role in contemporary philosophy of mind, psychology, cognitive science, and neuroscience, it is rarely defined. Allow me to illustrate the situation. In the field of psychology, the online *APA Dictionary of Psychology* says that awareness is “perception or knowledge of something”, and that “it is possible to be aware of something without being explicitly conscious of it”. This definition is clearly about transitive awareness, but it only allows us to infer that the terms ‘awareness’ and ‘consciousness’ are not coextensive. In the field of cognitive science, ‘awareness’ has no entry neither in the *MIT Encyclopedia of the Cognitive Sciences* (Wilson & Keil, 1999), nor in the *Dictionary of Cognitive Science* (Houdé, 2004). Finally, in the field of philosophy, the term has no entry neither in the *Stanford Encyclopedia of Philosophy*, nor in *The Oxford Dictionary of Philosophy* (Blackburn, 1996). Interestingly, the term has no entry in *The Oxford Companion to Consciousness* (Bayne et al., 2009), yet ‘awareness’ and its cognates appear more than 1,000 times throughout the whole volume.

Let us review how theorists use ‘awareness’ and ‘consciousness’ to determine whether there are enough reasons to support my claim that *consciousness* is not identical to *awareness*. In 1906, Dewey published a brief analysis of six different uses of the terms ‘consciousness’ and ‘conscious’. Two of those uses are particularly interesting. The fourth use says that: “‘Conscious’ means aware: ‘consciousness,’ the state of being aware. This is a wide, colorless use; there is no discrimination nor implication as to contents, as to what

there is awareness of, –whether mental or physical, personal or impersonal, etc.” (Dewey, 1906, p. 40). According to this “colorless” use, ‘consciousness’ points to intransitive awareness. Thus, since this use makes no reference to the contents of conscious states, i.e., to the fact that they are often about something, ‘consciousness’ names the general state of being awake.

The former use of ‘consciousness’ broadly coincides with our everyday use of the term ‘awareness’, which would explain why “to be conscious” and “to be aware” are often equivalent. Clearly, this use of ‘consciousness’ and ‘awareness’ is restricted to their intransitive senses. Nevertheless, although this use might support the idea that a system must be aware to have conscious states, i.e., that consciousness necessitates awareness, we still do not know why awareness does not necessitate consciousness. In consequence, it appears that we must look elsewhere for the difference between *consciousness* and *awareness*.

Dewey’s fifth use of ‘consciousness’ seems to point in the right direction. In its “distinctively philosophical use”, says Dewey, ‘consciousness’ “bring[s] out the difference between thoughts, [feelings, volitions] etc., characterizing the peculiar quality of a specific being or agent, and something which in general lies back of and conditions all such thoughts. Consciousness is now one with mind, or soul, or subject, as an underlying condition hypostasized into a substance” (Dewey, 1906, p. 40). While most contemporary theorists openly reject substance dualism (Descartes, 1984), they often characterize consciousness as, for instance, a concrete natural phenomenon (Chalmers, 1996, p. 128; G. Strawson, 2017, p. 84), a biological phenomenon (Searle, 2017, p. 331), or as a fundamental property (Goff, 2017, p. 3).

One of the advantages of “hypostasizing” consciousness is that it becomes a legitimate object of philosophical and scientific study. However, this strategy also opens the debate about the place of consciousness in reality. Consider that, although most contemporary theorists agree that “[c]onsciousness is a complex feature of the world” (Van Gulick, 2018), some have argued that *consciousness* is an ill-conceived concept (Churchland, 1995, p. 189), or that ‘consciousness’ refers to an illusory feature of experience (Dennett, 1991, p. 23). In particular, the claim that consciousness exists leads to the idea that *consciousness* refers to something that characterizes a special kind of entity called *subject* or *mind*.

Beyond the metaphysical implications of the views above, it should be noted that Dewey’s fifth use of ‘consciousness’ not only points to transitive consciousness, insofar as

it makes reference to the ability to entertain conscious states, like thoughts, feelings, volitions, etc., but also to the fact that those conscious states must be states of a system. While the previous “colorless” use of the term only referred to a kind of general state, this use refers to conscious states as internal states that characterize “the peculiar quality of a specific being or agent”. This is precisely the theoretically relevant sense of ‘consciousness’: when theorists use the term, they intend to call our attention to the fact that conscious states are experiences. It seems that we have found a good reason to claim that *consciousness* is not identical to *awareness*: to be aware of something is not the same as experiencing it.

The above are not the only technical uses of ‘consciousness’. For instance, a different way to describe our use of the term is by focusing on the kinds of systems capable of having conscious states. This aspect of the study of consciousness is called *creature consciousness* and is usually subdivided into animal consciousness and artificial consciousness, or machine consciousness. However, since I take all conscious things to be systems, it does not matter whether some of them are animal, artificial, celestial or imaginary, because a proper characterization of *consciousness* should be equally applicable to all conscious systems, as long as we take into account the relevant differences between the mechanisms that enable their conscious states. Consider that, while we know that dogs do not perceive colors exactly as we do, this does not necessarily mean that they are unable to entertain visual experiences. If it turns out that dogs are unable to entertain visual experiences, the reason cannot be that they perceive color differently.

Another technical use of ‘consciousness’ makes reference to the idea that conscious systems are cognitive systems. Accordingly, the system’s ontological status is superfluous, for conscious states could be realized by many ontologically different kinds of cognitive systems (Putnam, 1967). A concern about this use has to do with whether consciousness is binary or not, i.e., is consciousness an all-or-nothing state or does it come in degrees? Some have argued that it is not binary (Dennett, 1995; Lycan, 1996), while others have claimed that it is (Searle, 1994). Since this research is about *consciousness*, I do not need to address this debate.

We now have a way to explicate the difference between technical uses of ‘consciousness’ and ‘awareness’: while ‘awareness’ points to a certain cognitive ability, ‘consciousness’ points to experience. Hence, I take *awareness* to refer to the cognitive capacity to attend, discriminate, abstract, synthesize, and recall whatever a system’s perceptual apparatus is capable of encoding and processing. Briefly put, *awareness* points

to the functional basis of consciousness. Notice that, in contrast to Chalmers's (1996, p. 28) information-based definition of *awareness*, my characterization of the concept is given in cognitive terms. Conceptually speaking, a system could have awareness without having the capacity to encode information as we understand it, and not everything capable of encoding information is necessarily aware of it. Hence, despite Chalmers's reasoning (1996, pp. 293–297), a thermostat may have the capacity to encode and process information, but it hardly qualifies as something that experiences temperature changes.

In general, concepts can be characterized in different ways, for instance, in terms of the concept of its alleged fundamental property, by reducing it to another concept, or in relation to other concepts, i.e., by means of an inter-definition; this last approach is the one that I will support in the next chapter. It appears that we understand *consciousness* in different ways, i.e., according to different practices of conceptualization, even if we do not know exactly what consciousness is. Arguably, *consciousness* is not the only concept that can be characterized in different ways. Consider that, throughout the history of philosophy, *meaning* has been characterized in several incompatible ways, though, presumably, all of its characterizations point to the same phenomenon. In the following section, I will present three ways to characterize consciousness that can be drawn from different contemporary theories. With this, I intend to identify the conceptual background of those theories to determine the foundations of our understanding of *consciousness*. Hence, what follows is not an analysis of contemporary theories of consciousness, but a speculation about how theorists tend to define *consciousness*.

1.3. Three approaches to consciousness

Arguably, every theory has a conceptual background that determines the way in which its subject-matter is addressed. By 'conceptual background' I mean the characterization of a theory's subject-matter and how it relates to the rest of the concepts that are at the basis of the theory, thus forming a structure of interrelated concepts. In my view, part of what differentiates competing theories of consciousness is that they have different conceptual backgrounds, though they presumably share the same subject-matter⁶. In other words, while I grant that most contemporary theories of consciousness are, indeed, about consciousness, the fact that they adopt different theoretical strategies suggests that

⁶ All theories of consciousness are supposedly meant to explain the same phenomenon, so different characterizations of *consciousness* should at least have the same referent. Otherwise, they would be theories about different subject-matters, not competing theories of consciousness.

they explain different conceptually distinct aspects of consciousness. Thus, while many theorists have provided good reasons to endorse a theory over another, there does not seem to be a general concept of consciousness with which we could decide which theory best explains the nature of its referent.

The aim of this section is to reveal the foundations of our understanding of *consciousness* and, to that end, I will present three ways to characterize the concept that can be drawn from the apparent conceptual backgrounds of several contemporary theories of consciousness. I will argue that the following approaches to *consciousness* are practices of conceptualization that rely on the identification of different discernible features⁷ of conscious states, namely, intentionality, reflexivity and phenomenality, and that each of them assigns a higher theoretical weight to one of those features. Consequently, my intention is not to argue against any of those characterizations of *consciousness*, but to show that they address different aspects of consciousness.

First, allow me to discuss one particularly influential distinction, namely, that between *access consciousness* and *phenomenal consciousness*⁸ (Block, 1995). The first thing to note is that, since access consciousness is defined in terms of the availability of information for cognitive use (Block, 1995, p. 232), it is a functional feature. Thus, according to my previous analysis, ‘access consciousness’ does not point to consciousness, but to transitive awareness, insofar as a system could be able to process and use information without experiencing any of its internal states. This is not an argument against Block’s distinction because I think that it legitimately emphasizes the conceptual difference between the functional basis of consciousness, i.e., awareness, and consciousness. However, while Block presents his distinction as one between two concepts of consciousness, I believe that access consciousness is transitive awareness, not a concept of consciousness. Block’s second concept of consciousness, namely, phenomenal consciousness (Block, 1995, p. 230), does seem to point to at least part of the theoretically relevant sense of *consciousness*, insofar as it refers to the phenomenal character of experience⁹.

I will now speculate about the apparent conceptual backgrounds of different theories, to show that theorists mainly characterize *consciousness*: 1) in reference to the intentionality

⁷ I prefer to use the term ‘feature’, instead of ‘property’, because it has less metaphysical implications. Since this research is mainly focused on conceptual issues, I would like to leave open the debate on whether those features are actual properties of consciousness.

⁸ I thank Luke Roelofs and Sam Coleman for helping me improve my interpretation of Block’s distinction.

⁹ I will say more about the phenomenal character of experience later in this section.

of conscious states; 2) in terms of their reflexivity; or 3) in relation to their phenomenality. Briefly put, I will argue that *consciousness* has been explicated in terms of *intentionality*, *reflexivity* or *phenomenality*, which are concepts that point to different aspects of consciousness. Bear in mind that what follows is not an interpretation of how theorists explain consciousness, but an analysis of how they seem to characterize *consciousness*¹⁰.

1.3.1. Intentionality

One way to define *consciousness* is by expressing it in terms of the intentionality of conscious states, though defining *intentionality* has proven to be as difficult as defining *consciousness*¹¹. I will call this *the intentional approach*. By analyzing how theorists seem to explicate *consciousness* in terms of *intentionality*, I will show that the latter is a concept that points to a discernible feature of conscious states and that the intentional approach explains an aspect of consciousness.

Broadly speaking, the intentional approach is based on the thesis that “[m]any thoughts and experiences are ‘about’ or ‘of’ something” (Bayne et al., 2009, p. 382). Thus, whenever we experience something, or that something is the case, we entertain an intentional state, i.e., a kind of conscious state that has about-ness. To make things clear, I take ‘about-ness’ to point to the fact that many conscious states are about or of something¹², not to how their contents are represented. The sense of ‘intentionality’ that interests us here is the one that makes reference to

[...] the power of minds and mental states to be about, to represent, or to stand for, things, properties, and states of affairs. To say of an individual’s mental states that they have intentionality is to say that they are mental representations or that they have contents (Jacob, 2019).

Intentionality appears to be a distinctive feature of at least some conscious states, insofar as we seem to be able to identify and classify them according to their about-ness and their contents. Thus, the strategy of defining *consciousness* in terms of *intentionality* could be motivated by either the thesis that experiences are about or of something, or by the thesis

¹⁰ Since I will not assess any theory of consciousness, my analysis of how *consciousness* can be characterized may not coincide with how theorists actually do it. My goal is to reveal the foundations of our understanding of *consciousness* by analyzing the conceptual background that seems to support each view, not the views themselves.

¹¹ Some have even argued that it is impossible to understand one without the other (McGinn, 1991a, p. 78; Searle, 1992, p. 132).

¹² It should be noted that ‘intentionality’ “is a philosopher’s word” (Jacob, 2019), not to be confused with the ordinary sense of ‘intention’ as the purpose or goal of an action.

that experiences have contents. Notice that these are not equivalent theses¹³: an experience could be about something without having a specifiable content, while every experience that has a specifiable content is necessarily about something. Although we could endorse both theses¹⁴ and claim that the intentionality of conscious states consists in their about-ness and their having contents, the about-ness of a given state can still be conceptually detached from the state's content. Thus, it seems that about-ness is conceptually prior to the content of an intentional state, insofar the state must be about something to have a content. Furthermore, it is to be expected that, whenever an intentional state has a content, such content should be specifiable somehow¹⁵, which is not always the case.

In my view, the intentional approach exemplifies a practice of conceptualization where intentionality is taken as the fundamental property of consciousness. However, such practice is only possible once we have a relatively clear definition of *intentionality*. Thus, my suggestion is that the attribution of intentionality consists in the application of a concept. According to this practice of conceptualization, *intentionality* refers to the presentational character¹⁶ exhibited by some internal states. In consequence, every attribution of intentionality relies on the implementation of a convention regarding *intentionality* that ultimately depends on the conceptualization of the presentational character exhibited by some internal states.

Our understanding of *intentionality* seems to be mediated by our understanding of *about-ness*, so let us look closer into the latter. While some intentionalists, like Brentano, seem to suggest that about-ness is a feature or aspect of intentionality, others have claimed that “‘aboutness’ is the mere representation of some thing in words or thought, whether or not it exists” (Crane, 2013, p. 9). It seems to follow that every intentional state is a representational state, thus binding *about-ness* to the thesis that intentional states have specifiable contents. While it may seem that every human intentional state is also somehow representational, insofar as we are frequently able to represent the about-ness of our intentional states, it is not obvious that every intentional state is representational. In particular, it seems that the ability to represent exceeds the about-ness of intentional states.

¹³ Unless we take ‘about-ness’ to mean “having a content”.

¹⁴ Some theorists even treat ‘intentionality’ and ‘about-ness’ as synonyms (Bourget & Mendelovici, 2019).

¹⁵ For instance, as a proposition or a mental image.

¹⁶ The reason why I use the term ‘presentational’, instead ‘representational’, is that the former refers to the “particular form in which something is perceived by the mind”, while the latter refers to “[s]omething which stands for or denotes another symbolically”, as stated in their corresponding entries in the OED.

According to the OED, a representation is “[a]n image, concept, or thought in the mind, esp. as representing an object or state of affairs in the world”, i.e., representations are the contents of intentional states. Note that a representation can be formed verbally or by other cognitive means and that this is probably the reason why Crane suggests that about-ness consists in the representation of something “in words or thought”, where ‘thought’ points to non-verbal means¹⁷. However, we should not exclude the possibility of intentional states whose contents cannot be specified but that have about-ness¹⁸.

Conceptually speaking, we should distinguish *about-ness* from the fact that some intentional states have contents because contents must be represented by the system. This is the reason why the thesis that some internal states have about-ness is not equivalent to the thesis that they have contents. Thus, my suggestion is that *about-ness* points to the presentational character of some internal states. Nevertheless, some have insisted that intentional states are representational. Brentano, for instance, argued that “[e]very mental phenomenon includes something as object within itself” (Brentano, 2005, p. 68), where the objects that constitute the contents of intentional states are called *intentional objects*. However,

[i]t has sometimes been claimed [...] that some kinds of mental state are not intentional. The examples usually cited are somatic experiences, such as pain and pleasure, and moods, such as depression and elation. On the face of it, a toothache is not about anything, and nor is an orgasm. Likewise, depression, anxiety, and elation do not appear to be about anything in particular. Yet all these are clearly mental phenomena (Bayne et al., 2009, p. 383).

Brentano’s followers frequently reply that “somatic experiences are about something, after all: they are about physiological events in one’s body” (Bayne et al., 2009, p. 383). However, it seems that we are dealing with two senses of ‘about-ness’: on the one hand, the term points to the presentational character exhibited by some conscious states; on the other, it points to the representation of the intentionality of a state. Consider, for instance, the about-ness of a headache: while, in the second sense, it could be said that the headache is about an inflammation in my head; in the first sense, my experience of the headache is about its phenomenality, i.e., what it is like for me to experience the headache. Accordingly,

¹⁷ I thank Sam Coleman for helping me improve my interpretation of Crane’s characterization of about-ness.

¹⁸ Arguably, pre-linguistic children entertain states that exhibit a presentational character, i.e., that have about-ness, though they might not yet have the cognitive skills to represent what their states are about. Furthermore, though all unimpaired adult humans possibly share the same perceptual apparatus and the same representational skills, different cultures tend to represent similar perceptions differently (Miyamoto et al., 2006).

it seems that we are dealing with two different practices of conceptualization concerning *about-ness*, one directed at the conceptualization of the presentational character of an experience, the other at the representation of the intentionality of the experience.

The latter analysis of the intentionality of somatic experiences evinces a close link between *intentionality* and *phenomenality*, insofar as the about-ness of somatic experiences seem to point to their what-it's-like-ness. Additionally, since somatic experiences are actual states that conscious systems entertain under certain circumstances, their intentionality must also be linked to the system's ability to entertain states that exhibit reflexive characters, i.e., its ability to be aware of itself and its own mentality. Hence, there is also a link between *intentionality* and *reflexivity*. I will have more to say about these links in the third chapter, but it should be noted that the closer we look into *intentionality*, the closer we get to the other two discernible features of conscious states, namely, *reflexivity* and *phenomenality*.

So far, we have found a direct relation between *intentionality* and *reflexivity*, insofar as some intentional states require the system to be able to entertain states that exhibit reflexive characters. In consequence, it seems that at least some attributions of intentionality also involve the attribution of reflexivity, which suggests that *intentionality* is not independent from *reflexivity*. Thus, although intentionality can be conceptually distinguished from reflexivity, it does not follow that they must be independent properties of consciousness. This is the reason why I treat them as discernible features of conscious states.

We have also seen that the intentional approach is grounded on the thesis that all conscious states are intentional states because they are mainly identified by their about-ness. However, there are cases where the intentionality of conscious states involves non-intentional features, like the phenomenality and the reflexivity of moods, emotions, somatic experiences, etc. Hence, it seems that *intentionality* cannot be detached from *reflexivity* and *phenomenality*, which suggests that *consciousness* cannot be fully characterized in terms of *intentionality*, unless we make some further claims¹⁹. In any case, I have shown that a proper characterization of *intentionality* requires an explicit reference to *phenomenality* and *reflexivity*.

¹⁹ For instance, an intentionalist could argue that the phenomenality of an experience is part of its intentionality (Tye, 1995, p. 137) and that reflexivity is not a feature of the experience, but an instance of a principle that applies to consciousness in general (Roelofs, 2016, p. 3205).

My goal here is not to say that the intentional approach is flawed or misguided. On the contrary, I think that the intentional approach provides a good explanation of an aspect of consciousness. What I have tried to argue is that *consciousness* could only be characterized in terms of *intentionality* by taking into account its connection to *phenomenality* and *reflexivity*. Otherwise, we would not be able to account for conscious states that do not seem to exhibit a presentational character, like moods and emotions. For these reasons, I do not think that there is an intentionalist concept of consciousness, i.e., there is nothing like *intentional consciousness*. In my view, the intentional approach relies on a specific way to characterize *consciousness*, namely, according to a practice of conceptualization where intentionality is taken as the fundamental property of consciousness.

My claim is not that the intentional approach fails to capture *consciousness* and that, consequently, the theories that seem to endorse this approach are unable to explain consciousness. Instead, what I have tried to argue is that, while the intentional approach adopts a theoretical strategy where intentionality is posited as the fundamental property of consciousness, the way in which *consciousness* is characterized involves the concepts of the other two discernible features of conscious states, namely, *reflexivity* and *phenomenality*. Thus, although the intentional approach seems to be suitable for explaining an aspect of consciousness, it is not obvious that an explanation of the ultimate nature of consciousness could be given in purely intentional terms, i.e., independently of *phenomenality* and *reflexivity*.

1.3.2. Reflexivity

Another way to characterize *consciousness* is by referring to the reflexivity of conscious states. I will call this *the reflexive approach*. According to some theorists (Gennaro, 2012; Kriegel, 2009), conscious states are reflexive²⁰, insofar as they involve self-awareness or self-examination. Broadly speaking, the reflexive approach is grounded on the intuition that conscious systems have the ability to be aware of at least some of their internal states²¹. In particular, theorists tend to point out that “[h]uman beings are conscious not only of the world around them but also of themselves: their activities, their bodies, and their mental lives” (Smith, 2017). Thus, it would seem that conscious states always exhibit a

²⁰ It should be noted that, although ‘reflexivity’ is often related to ‘introspection’, ‘subjectivity’, and ‘perspectivity’, the terms are not coextensive.

²¹ I thank Sam Coleman for helping me clarify the foundations of the reflexive approach.

reflexive character. As in the case of *intentionality*, by analyzing how theorists seem to explicate *consciousness* in terms of *reflexivity*, I will show that the latter is a concept that points to a discernible feature of conscious states and that the reflexive approach explains an aspect of consciousness.

In my view, the reflexive approach exemplifies a practice of conceptualization where reflexivity is taken as the fundamental property of consciousness. However, such practice is only possible once we have a relatively clear definition of *reflexivity*. Thus, my suggestion is that the attribution of reflexivity consists in the application of a concept. According to this practice of conceptualization, *reflexivity* refers to the reflexive character exhibited by some internal states. Thus, every attribution of reflexivity relies on the implementation of a convention regarding *reflexivity* that ultimately depends on the conceptualization of the reflexive character exhibited by some internal states.

To say that conscious states are reflexive does not necessarily mean that they are self-referential. Actually, it could be argued that self-referentiality is not a property that is exclusive of conscious states²². In contrast, the advocates of the reflexive approach tend to argue that all conscious states exhibit a reflexive character, insofar as they involve the system's ability to be aware of at least some of its internal states. Thus, according to this approach, a system is conscious if it can be aware of some of its internal states, not if it entertains self-referential states.

In a sense, *reflexivity* is related to what some theorists call *mine-ness* because it points to the fact that some internal states are experienced by the system as its own states²³. Accordingly, one might argue that reflexivity is a lower-order²⁴ property of consciousness often called *pre-reflective self-consciousness* (Kriegel, 2009, pp. 176–181), *minimal self-awareness* (Zahavi, 2005), or *weak self-consciousness* (Flanagan, 1995, p. 193), and it could also be interpreted in terms of what Roelofs (2016, p. 3205) calls the *Ownership Principle*. What matters is that *reflexivity* is a concept that points to the reflexive character exhibited by some internal states. For instance, in seeing an apple, I am conscious of my perception of

²² For instance, it is often a property of linguistic expressions, works of art, mathematical and logical formulas, and computer programs.

²³ Nevertheless, *reflexivity* and *mine-ness* are not the same concept. While the former points to a feature of at least some conscious states, the latter points to how the system experiences such feature. Thus, *mine-ness* seems to result from a relation between *reflexivity* and *phenomenality*.

²⁴ It might seem that my discussion of the discernible features of conscious states is actually about lower-order properties of consciousness. However, since I am mainly interested in conceptual aspects, the discernible features of consciousness may not resemble the actual properties of consciousness.

it, even if I cannot represent myself as such-and-such. Furthermore, consider that newborns might not be conscious of themselves as selves, though they do seem to entertain states that exhibit reflexive characters (Flanagan, 1995, p. 193).

Theorists have proposed different ways to identify the reflexive character of conscious states. For instance, some have pointed out that we, and possibly many other kinds of systems, are capable of bodily awareness: “not only do we perceive it [our body] through external senses, but we have also an internal access to it through bodily sensations” (de Vignemont, 2018). Thus, *bodily awareness* points to the relation we have with our bodies by means of our proprioceptive skills: we control, monitor and feel our own bodies “from the inside”. In my view, every experience of our own body is necessarily reflexive, insofar as it implies having awareness of the experience, i.e., there cannot be an experience of our body of which we are not aware because reflexivity is part of the experience itself²⁵. Otherwise, we would have to claim that we could experience, say, the position of our body, without being aware of it.

In any case, *bodily awareness* points to a kind of reflexivity that only applies to embodied systems and, conceptually speaking, there could be disembodied conscious systems. Furthermore, that conscious embodied systems tend to have bodily awareness does not mean that they are aware of every state of their body: I am certainly not aware of the level of serotonin in my body, though I can be aware of my mood²⁶.

The interpretation of *bodily awareness* is controversial and there is no general agreement about how to explain its referent. Notice, for instance, that bodily awareness is not the same as having self-monitoring mechanisms. Our bodies have several self-monitoring mechanisms, like temperature and heart rate regulation, of which we cannot be aware. Analogously, computers are often programmed to monitor their internal states, but they do not seem to be able to experience those states. Consequently, *bodily awareness* should not be reduced to the purely mechanical processes of regulation and monitorization that take place in a system’s body²⁷.

²⁵ That we are not always able to conceptualize the reflexivity of states of bodily awareness has to do with the fact that sometimes we are not capable of identifying their reflexive characters, not with their absence.

²⁶ That the level of serotonin in my body is related to my mood is not something I can experience because it is not part of my ability to entertain states of bodily awareness. Put briefly, I can experience mood variations, not changes in the level of serotonin in my body.

²⁷ Note that states of bodily awareness cannot be explained in terms of their physiological basis, i.e., in terms of *awareness*, because they exhibit a reflexive character that cannot be reduced to the self-monitoring mechanisms that enable them.

Moreover, that sometimes we are unable to describe the reflexive character of a state of bodily awareness, as when we have a sensation that we cannot describe, only means that some reflexive characters cannot be described in terms of *intentionality*. More importantly, since only embodied systems can entertain states of bodily awareness, a characterization of *consciousness* in terms of this kind of reflexivity would exclude the possibility of there being disembodied conscious systems. Besides, it seems clear that we entertain many conscious states that are not states of bodily awareness, though they often exhibit reflexive characters. Since *bodily awareness* points to the reflexive character of a subset of our conscious states, namely, those of some of the states of our body, *consciousness* cannot be fully characterized in terms of this kind of reflexivity.

Another way to characterize the reflexivity of conscious states is by means of what Bermúdez (1998), Peacocke (2014) and Zahavi (2005), among others, call *self-consciousness* or *self-awareness*²⁸. Broadly speaking, the idea is that, “[i]f one is capable of self-conscious thought, [...] one must be able to think in such a way that it is manifest to one that it is oneself about whom one is thinking” (Smith, 2017). Notice that, in contrast to bodily awareness, self-conscious states are reflexive in the sense that they imply having awareness of our own mentality, not about the states of our body. While *bodily awareness* specifically refers to embodied systems, *self-consciousness* points to the system’s mentality, which means that there could be disembodied self-conscious systems.

Self-conscious states involve the system’s ability to be aware of itself and of its own mental states. Thus, a system that entertains self-conscious states is a system capable of experiencing its own mentality. This characteristic of self-conscious states suggests that *self-consciousness* can be linked to *intentionality*, inasmuch as self-conscious states are states about the system’s internal states²⁹. However, the reflexivity of self-conscious states does not necessarily imply that they also have a distinctive intentionality. From the point of view of the reflexive approach, one could argue that the reflexivity of a self-conscious state does not mean that the state has a content or even that it is about something in the relevant sense: I can be conscious, or aware³⁰, of myself without being, at the same time, aware that my consciousness of myself is about myself or that its content is a representation of myself.

²⁸ It is noteworthy that ‘self-awareness’ and ‘self-consciousness’ are equivalent because they both point to the intransitive senses of ‘awareness’ and ‘consciousness’.

²⁹ I will have more to say about the link between *self-consciousness* and *intentionality* in the third chapter.

³⁰ Recall that ‘aware’ and ‘conscious’ are equivalent in this context because they point to the transitive senses of ‘awareness’ and ‘consciousness’.

Briefly put, to be aware of myself and to be aware of my awareness are two different kinds of conscious states.

It could also be argued that self-conscious states exhibit phenomenal characters because, after all, they are experienced by the system. Thus, it would seem that *self-consciousness* is linked to *phenomenality*³¹. However, it is not clear how the phenomenal characters of self-conscious states could be described in terms of *reflexivity*. Moreover, at least from the conceptual point of view, there might be self-conscious states that do not exhibit phenomenal characters. Hence, it seems that, while *self-consciousness* points to the reflexivity of many of our conscious states, it cannot account for their other features, which means that *consciousness* cannot be fully characterized in terms of *self-consciousness*.

A further way to characterize the reflexivity of conscious states is by means of *self-knowledge*. According to some theorists (Descartes, 1984, p. 20; Locke, 1975, p. 335; Nozick, 1981), we are not only self-conscious, we can also adopt several propositional attitudes towards our mentality. Since Descartes (1984), many theorists have claimed that our self-knowledge is a) infallible, because it seems impossible for us to have false beliefs about being in a certain state; and b) omniscient, because being in a certain conscious state is sufficient for knowing that we are in that state. However, *self-knowledge* clearly refers to higher-order intentional states, insofar as knowing and believing are propositional attitudes that not only require the system to be self-conscious, but also to be able to entertain intentional states that target its own conscious states. Thus defined, although *self-knowledge* seems to link *reflexivity* and *intentionality*, it points to a reduced set of conscious states. Consider that, while my perception of an apple probably exhibits the kind of reflexivity to which *bodily awareness* points, it is not a state of self-knowledge, insofar as I do not need to know that I am perceiving an apple to be aware of my perception of it.

Some theorists (Fasching, 2009; Slors & Jongepier, 2014; Zahavi, 2000, 2005) have claimed that the reflexivity of conscious states is not an extra ingredient, nor a property of consciousness, because most, if not all, conscious states exhibit a reflexive character; conceptualizable or not. For this reason, one could claim that reflexivity is not a feature of conscious states, but a condition of consciousness. However, it should be noted that not all conscious states can be described in terms of *reflexivity* because a description of the reflexive character of a state is not a description of its intentionality or its phenomenality.

³¹ I will have more to say about the link between *self-consciousness* and *phenomenality* in the third chapter.

The latter suggests that *consciousness* cannot be fully characterized in terms of *reflexivity* because it is closely connected to *intentionality* and *phenomenality*. Moreover, it seems that the reflexivity of many conscious states is somehow modified by their intentionality and their phenomenality, as well as it also seems to modify them. Thus, I do not think that there is a reflexive concept of consciousness, i.e., there is nothing like *reflexive consciousness*. In my view, the reflexive approach relies on a specific way to characterize *consciousness*, namely, according to a practice of conceptualization that posits reflexivity as the fundamental property of consciousness.

My claim is not that the reflexive approach fails to capture *consciousness* and that, consequently, the theories that seem to endorse this approach are unable to explain consciousness. Instead, what I have tried to argue is that, while the reflexive approach adopts a theoretical strategy where reflexivity is posited as the fundamental property of consciousness, the way in which *consciousness* is characterized involves the concepts of the other two discernible features of conscious states, namely, *intentionality* and *phenomenality*. Thus, although the reflexive approach seems to be suitable for explaining an aspect of consciousness, it is not obvious that an explanation of the ultimate nature of consciousness could be given in purely reflexive terms, i.e., independently of *intentionality* and *reflexivity*.

1.3.3. Phenomenality

A third way to characterize *consciousness* is by referring to the fact that experience is like something (Nagel, 1974), i.e., that there is something it is like to entertain a conscious state. In other words, the idea is that conscious states exhibit a phenomenal character³², insofar as they seem to have phenomenal properties³³. I will call this *the phenomenal approach*. Phenomenalists (Block, 1995; Chalmers, 1995, 1996; Jackson, 1986; Levine, 1983; Nagel, 1974) often argue that a proper explanation of consciousness must account for its phenomenality. As in the cases of *intentionality* and *reflexivity*, by analyzing how theorists seem to explicate *consciousness* in terms of *phenomenality*, I will show that the latter is a concept that points to a discernible feature of conscious states and that the phenomenal approach explains an aspect of consciousness.

The phenomenalist approach is grounded on the thesis that there is “always ‘something it is like’ to be in a given conscious state –something it’s like for one who is in

³² Often referred to as “the subjective quality of experience” (Chalmers, 1996, p. 4).

³³ Also called *qualia*.

that state— and what it’s like for you to be in a state is what makes it a conscious state of the kind it is” (Siewert, 2016). In consequence, according to this approach, *consciousness* must be described in terms of *phenomenality* or, more precisely, *consciousness* and *phenomenal consciousness* are coextensive.

In my view, the phenomenalist approach exemplifies a practice of conceptualization where phenomenality is taken as the fundamental property of consciousness. However, such practice is only possible once we have a relatively clear definition of *phenomenality*. Thus, my suggestion is that the attribution of phenomenality consists in the application of a concept. According to this practice of conceptualization, *phenomenality* refers to the phenomenal character exhibited by some internal states. Thus, every attribution of phenomenality relies on the implementation of a convention regarding *phenomenality* that ultimately depends on the identification of the phenomenal character exhibited by some internal states.

Broadly speaking, *phenomenality* points to the what-it’s-like-ness of experience which, according to my view, is a conceptualization of the phenomenal character of conscious states. Hence, we must start by exploring *what-it’s-like-ness*. According to many theorists, the what-it’s-like-ness of conscious states is constituted by phenomenal features that are irreducible to their cognitive/psychological and functional features. Moreover, it seems that specific kinds of experiences, like perceptions, have specific phenomenal features. For instance, we can focus on our perception of a tomato and describe its color in terms of what it is like for us to perceive it. Furthermore, according to some theorists (G. Strawson, 2004), thoughts and beliefs also have distinguishable phenomenal features: the what-it’s-like-ness of thinking about global warming seems to be characteristically different from the perception of a tomato. Hence, one could argue that the what-it’s-like-ness of experience is at least partially determined by the modality³⁴ of the experience which means that each modality is accompanied by different phenomenal features.

A further characteristic of the what-it’s-like-ness of experience is that it is context-sensitive (Roelofs, 2014, p. 64), in the sense that the phenomenal features of an experience depend on environmental, physiological and cultural factors: perceptions are certainly constraint by external conditions, a system’s perceptual and cognitive mechanisms influence how the system experiences under different circumstances, and the cultural environment of

³⁴ These modalities are the different means through which a system could entertain conscious states, namely, sensory perception, proprioception, thought, belief, etc.

a system shapes the way in which the system interprets its experiences. Besides these three elements, there is the cognitive/psychological element that concerns the overall mentality of the system, like its previous experiences and its general state of consciousness³⁵. Thus, the what-it's-like-ness of, say, my belief in global warming results from the integration of these four elements.

Another way to characterize *what-it's-like-ness* is by means of *subjectivity*. It is often said that there is a certain epistemic asymmetry between our own conscious states and those of others: “Facts about conscious experience can be at best incompletely understood from an outside third person point of view” (Van Gulick, 2018). We might be able to somehow share with others some phenomenal features of our conscious states (Roelofs, 2016, pp. 3217–3218; Velmans, 2009, p. 212), but this phenomenon cannot be explained by means of the standard third-person methods of contemporary science (Velmans, 2009, pp. 219–221). Thus, the subjectivity of conscious states consists in the fact that their what-it's-like-ness could only be indirectly accessed by having a conscious state of the same sort (Velmans, 2009, pp. 212–213). Furthermore, some have claimed that the subjectivity of our conscious states can only be understood “by creatures like us” (Nagel, 1974, p. 440), which implies that the phenomenality of conscious states is relative to each kind of conscious system.

It is noteworthy that *subjectivity* does not point solely to the phenomenality of experience, but also to its reflexivity and its intentionality³⁶: arguably the phenomenality of my perception of a tomato is not identical to yours, though their phenomenal features might be the similar³⁷, and it could also happen that the phenomenality of my belief in global warming is different from yours, though both of our beliefs would be about the same thing³⁸. Thus, it would seem that an explanation of the phenomenality of a given experience cannot be detached from its intentionality or from the fact that it is experienced by some system. Moreover, it is plausible that the phenomenality of conscious states is relative to the kind of system that entertains them and, perhaps, to each particular system. Indeed, pain may feel differently for different kinds of systems, but it must be painful. If squirrels are capable of

³⁵ Namely, the system's intransitive consciousness.

³⁶ I will say more about this in the third chapter.

³⁷ For instance, my cultural environment, which influences the reflexivity of my conscious states, could modify the what-it's-like-ness of my perception, thus making it different from the what-it's-like-ness of your perception.

³⁸ Variations in the way in which the content of a belief is represented often lead to variations in the phenomenality of the belief. Suppose that we both believe in global warming, but we represent it differently. For instance, one of us could further believe that we are at least partially responsible for it, while not the other. Although we both believe in global warming, the variations in our representations of it will probably generate a difference in the phenomenality of our beliefs.

feeling pain, it must certainly be painful to them, even if we cannot access the particular what-it's-like-ness of their painful experiences. Consequently, a characterization of *phenomenality* in terms of *subjectivity* would require making explicit its connections to *intentionality* and *reflexivity* because it seems that the what-it's-like-ness of an experience can be modified by the experience's intentionality and reflexivity.

A further way to characterize *what-it's-like-ness* is through the idea that consciousness seems to be a unified phenomenon, in the sense that the features of conscious states seem to be integrated. For instance, it is plausible that the intentionality of our perceptions results from the integration of different aspects of our perceptual field: "We can say that two states of consciousness are objectually unified when they are directed at the same object" (Chalmers & Bayne, 2003, p. 24). Analogously, our conscious states seem to be phenomenally unified: "[i]t is difficult or impossible to imagine a subject having two phenomenal states simultaneously, without there being a conjoint phenomenology for both states" (Chalmers & Bayne, 2003, p. 37).

According to the thesis that consciousness is a unified phenomenon, my perception of a tomato is not only intentionally unified, but also phenomenally, insofar as the what-it's-like-ness of my perception is not just a composition of independent phenomenal features. Analogously, when we are at a music concert, although we experience several perceptions that come from different senses, they all seem to be unified into a single representational/phenomenal field. This characterization of *what-it's-like-ness* evinces a close link between *phenomenality* and *intentionality*, insofar as the unity of the what-it's-like-ness of an experience seems to be relative to the unity of its intentionality, and vice versa.

My analysis of the above ways to define *what-it's-like-ness* suggests that *consciousness* cannot be fully characterized in terms of *phenomenality* because the phenomenality of experience is often linked to its intentionality and its reflexivity. As a matter of fact, it would seem that the phenomenality of many conscious states depends or is somehow modified by their intentionality and their reflexivity. Accordingly, I do not think that there is a phenomenal concept of consciousness, i.e., there is nothing like *phenomenal consciousness*. In my view, the phenomenal approach relies on a specific way to characterize *consciousness*, namely, according to a practice of conceptualization where phenomenality is posited as the fundamental property of consciousness.

My claim is not that the phenomenal approach fails to capture *consciousness* and that, consequently, the theories that seem to endorse this approach are unable to explain consciousness. Instead, what I have tried to argue is that, while the phenomenal approach adopts a theoretical strategy where phenomenality is posited as the fundamental property of consciousness, the way in which *consciousness* is characterized involves the concepts of the other two discernible features of conscious states, namely, *reflexivity* and *intentionality*. Thus, although the phenomenal approach seems to be suitable for explaining an aspect of consciousness, it is not obvious that an explanation of the ultimate nature of consciousness could be given in purely phenomenal terms, i.e., independently of *intentionality* and *reflexivity*.

In this chapter, I have argued that intentionality, reflexivity, and phenomenality are discernible features of conscious states. According to my analysis of how *intentionality*, *reflexivity*, and *phenomenality* are often defined, they can be employed to characterize *consciousness* in different ways, by means of adopting different practices of conceptualization. However, it seems that our characterizations of *intentionality*, *reflexivity*, and *phenomenality* involve their interrelations, which is the reason why I treated them as features, instead of properties. What I have shown in this section is that *consciousness* can be defined according to different strategies and that each one of them addresses different aspects of consciousness, not that they are fundamentally mistaken. Hence, I suspect that the adoption of any of the approaches above is a pragmatic choice, in Dennett's sense (1971), while it is not obvious whether any of them provides the ultimate characterization of our subject-matter.

The aim of this chapter was to show that *consciousness* can be explicated in different ways, which opposes the thesis that there are many concepts of consciousness. Indeed, our uses of 'consciousness' and our understanding of *consciousness* may be vague, but that does not mean that different characterizations point to different referents. Regardless of what consciousness is, *consciousness* must be a single concept that points to a single referent. In the next chapter, I will introduce an interpretation of *consciousness* that takes *intentionality*, *reflexivity* and *phenomenality* as its components, thus unifying the different characterizations analyzed in this section in terms of a single conceptual space.

Chapter 2: The Concept of Consciousness

Bart: Um, Dad?

Homer: Yeah?

Bart: What is the mind? Is it just a system of impulses, or is it something tangible?

Homer: Relax! What is mind? No matter. What is matter? Never mind.

The Simpsons, 'Good night', 1987

In the previous chapter, I have analyzed our uses of the terms 'awareness' and 'consciousness', as well as three possible ways to characterize *consciousness* that can be drawn from different theories of consciousness. According to that analysis, our characterization of *consciousness* is mediated by our characterization of *intentionality*, *reflexivity* and *phenomenality*, and the way in which we characterize *consciousness* depends on our theoretical interests. The goal of that chapter was to show that there are different ways to characterize a single concept according to different practices of conceptualization, not independent concepts of consciousness.

In this chapter, I will explore the hypothesis that the competing approaches to *consciousness* that I presented in the previous chapter can be reconciled by highlighting the interdependence between *intentionality*, *reflexivity* and *phenomenality*. Thus, the goal of this chapter is to provide an interpretation of *consciousness* that is general enough to fit those three seemingly independent ways to characterize the concept, which is why I believe that my interpretation of the concept is metaphysically compatible with most contemporary theories of consciousness and with different views on the ultimately correct analysis of *consciousness*. While theorists normally take intentionality, reflexivity and/or

phenomenality to be properties of consciousness, I will argue that their concepts are components of *consciousness* and that their interdependence constitutes a single conceptual space. Hence, it might turn out that my interpretation of *consciousness* does not resemble the nature of its referent, though it is meant to correspond to it. Accordingly, my goal is not to present a theory of consciousness, but a conceptual background on which an alternative kind of theory could be based.

2.1. An interpretation of the concept of consciousness

One way to introduce my interpretation of *consciousness* is by explaining why I think that intentionality, reflexivity and phenomenality are the discernible features of conscious states, instead of others. Surely, conscious states appear to have other features that are not reducible to the former, so it might seem that my interpretation of *consciousness* will not be able to fully capture the concept³⁹.

I believe that all conscious states can be described in terms of relations between *intentionality*, *reflexivity* and *phenomenality*. Although I will elaborate on this idea in the next chapter, a few examples might illustrate my view. Self-consciousness, understood as the capacity to have states about what it is like to be a particular system, can be interpreted as a relation between the reflexive and the phenomenal characters of some internal states. Hence, *self-consciousness* could be characterized in terms of *reflexivity* and *phenomenality*⁴⁰. Self-knowledge, understood as the capacity to hold propositional attitudes towards our mentality, can be interpreted as a relation between the reflexive and the intentional characters of some internal states. Thus, *self-knowledge* could be characterized in terms of *reflexivity* and *intentionality*⁴¹. Subjectivity, understood as the particular way in which some states are like for a particular system, can be interpreted as a combination of all three features, insofar

³⁹ Keep in mind that I did not claim that the three approaches of *consciousness* that I discussed in the previous chapter are mistaken. What I argued is that they address conceptually distinct aspect of consciousness because they are motivated by different theoretical interest or practices of conceptualization. Put briefly, they all characterize the same concept differently, namely, by privileging one of the discernible features over the others.

⁴⁰ *Intentionality* is also involved, insofar as we are dealing with states that exhibit presentational characters. However, it should be noted that the about-ness of self-conscious states often points to their phenomenality, so their intentionality originates from the relation between their reflexive and their phenomenal characters.

⁴¹ Although this kind of conscious state could also exhibit phenomenal characters, their phenomenality is not what characterizes them, so *phenomenality* is not essential to understanding *self-knowledge*. Consider that the phenomenal character of knowing that I had a pizza last night does not seem to be relevant to describe it as a state of self-knowledge. Furthermore, notice that *self-knowledge* is not equivalent to *self-consciousness* because a self-conscious system does not need to be able to hold propositional attitudes towards its own mentality to experience what it is like to be the system it is.

as subjective states are about what the state is like for a particular system. Therefore, *subjectivity* could be characterized in terms of *intentionality*, *reflexivity* and *phenomenality*⁴².

Another way to introduce my interpretation of *consciousness* is by distinguishing the terminology usually employed to study the nature of consciousness from the one that I will adopt to characterize its concept. Theorists normally differentiate subjects, or referents, from their properties. Accordingly, as I have shown in the previous chapter, they often argue that consciousness, which is a subject or referent, can be explained by characterizing its alleged properties, namely, intentionality, reflexivity and phenomenality. Of course, theorists do not universally agree on whether the latter are all properties of consciousness, but their approaches follow a similar methodology: to explain consciousness, we must identify and characterize its fundamental property or properties.

Since I do not know what consciousness is, or at least this is the initial assumption of my view, I only take for granted that consciousness is something, it is our subject-matter. Accordingly, I also do not know whether intentionality, reflexivity and phenomenality are actual properties of consciousness, though it seems that they are at least features of conscious states⁴³ and, for that reason, I call them *discernible features of conscious states*. The idea is that, regardless of whether they are actual properties of consciousness, they are conceptually distinct aspects of it. Now, in the previous chapter, I showed that *consciousness* can be characterized by means of the concepts of the discernible features, namely, *intentionality*, *reflexivity* or *phenomenality*. Because these are concepts, not properties, I call them *components of the concept of consciousness*. Notice that intentionality and reflexivity could be features of other things, besides conscious states, though not phenomenality because the latter is generally characterized in relation to consciousness or even identified with it. In any case, it seems that the concepts of the discernible features are somehow intertwined with *consciousness*, which is evinced by the ways in which theorists seem to characterize *consciousness*. Finally, it should be noted that, just as any other subject-matter, consciousness, and especially conscious states, presents itself in specific ways. I call those ways or modes of presentation *characters*.

I will focus on the characters of conscious states in the next chapter, but it is noteworthy that they are at the base of our conceptualizations of the discernible features, i.e.,

⁴² Notice that *subjectivity* is not equivalent to *self-consciousness* because the former points to the ability to entertain states about what a particular conscious state is like for a particular system, not just about what it is like to be that system.

⁴³ Whatever they are from the metaphysical point of view.

the components of *consciousness*, in the sense that we identify and describe conscious states depending on how they present themselves. In other words, the characters of conscious states are the ways in which consciousness manifests itself, and they are conceptualized as *intentionality*, *reflexivity* and *phenomenality*. Thus, in contrast to the usual terminology, based on subjects and properties, I will introduce my interpretation of *consciousness* in terms of features, components and characters.

In a sense, my interpretation of *consciousness* is based on a circular reasoning⁴⁴ because I will claim that our understanding of *consciousness* is mediated by our understanding of *intentionality*, *reflexivity* and *phenomenality*, while our understanding of those concepts relies on our understanding of *consciousness*. However, in my view, *consciousness*, *intentionality*, *reflexivity* and *phenomenality* are interdependent concepts, which means that their characterization results from how we conceive their relations, not from the primacy of any of them⁴⁵. Hopefully, the circularity of my proposal will become sufficiently “wide, revealing, and illuminating” (Strawson, 1992, pp. 19–20) to undermine the objection.

In my view, *consciousness* is constituted by *intentionality*, *reflexivity* and *phenomenality*, thus creating a sort of three-dimensional conceptual space⁴⁶ where each of the concepts of the discernible features of conscious states is a dimension of the space. Put briefly, *intentionality*, *reflexivity* and *phenomenality* are the components of the conceptual space of *consciousness*. Notice that, since I am dealing with concepts, not with their referents, what my interpretation implies is that the components of *consciousness* can be rationally distinguished from each other⁴⁷, though they are interrelated with respect to the conceptual space of *consciousness*. Thus, the conceptual space of *consciousness* results from the interrelations between its components. In contrast to the approaches that I have analyzed in the previous chapter, this interpretation suggests that there is one concept of consciousness and that its conceptual space can be modified according to our theoretical interests, as I will

⁴⁴ I thank Glorymar Hernandez for calling my attention to this issue.

⁴⁵ This is another reason why I believe that my view is compatible with different views regarding the ultimately correct analysis of *consciousness*. Later in this chapter, I will argue that the conceptual space of *consciousness* can be customized to form models of consciousness and such models will correspond to different interpretations of the nature of consciousness.

⁴⁶ My interpretation of *consciousness* could be put in terms of what Márquez Velasco (2013, p.101) calls a *spatio-conceptual coordinate*, insofar as it is meant to make explicit “a series of potential connections” (Márquez Velasco, 2013, p. 100) between *consciousness* and its components. It could also be put in terms of Gärdenfors’s (2000) *Conceptual Space Theory*, where *consciousness* is formed by three integral quality dimensions.

⁴⁷ This is the reason why I call the features of conscious states *discernible*.

argue later in this chapter. Accordingly, there are different configurations of the same conceptual space, not different concepts of consciousness.

There are two reasons why I have decided to represent *consciousness* as a three-dimensional conceptual space: a) it provides a conceptual background with which different approaches to *consciousness* can be assessed; and b) it motivates an alternative way of theorizing about consciousness. Regarding a), it should be noted that the three approaches to *consciousness* that I have analyzed in the previous chapter can be seen as different configurations of the conceptual space that assign a higher theoretical weight to a subset of the components of *consciousness*. As for b), I can advance that it will be the target of the next chapter, where I will introduce a way to describe conscious states that solely relies on my interpretation of *consciousness*.

The foundation of my interpretation of *consciousness* is the thesis that paradigmatic conscious states exhibit either of three identifiable characters, namely, the presentational, the reflexive and the phenomenal. My previous analysis of different approaches to *consciousness* showed that intentionality, reflexivity and phenomenality are discernible features of conscious states, and that some of those states can only be identified by means of conceptualizing the relations between the characters they exhibit, which is why I take the concepts of those features to be interrelated or interdependent. Hence, my proposal is that, although *consciousness* can be characterized according to different theoretical interests, its conceptual space results from the interrelations between its components. My interpretation of *consciousness* is meant to highlight that the connections between *intentionality*, *reflexivity* and *phenomenality* can be interpreted in terms of a single conceptual space. For illustrative purposes, the conceptual space of *consciousness* can be thus represented:

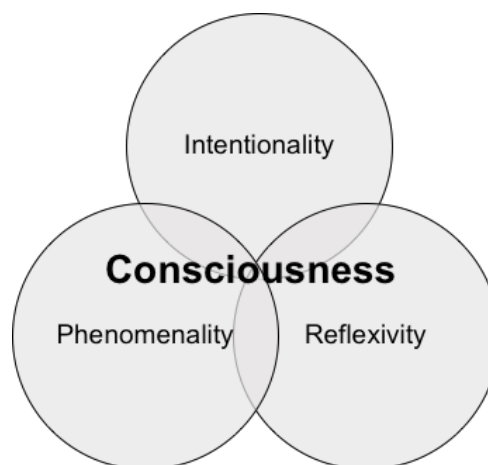


Figure 1: The Conceptual Space of Consciousness

The conceptual space of *consciousness* is a heuristic device that says nothing about the ultimate nature of consciousness because I have built it from the concepts of the discernible features of conscious states, not from examining the alleged properties of consciousness⁴⁸. More importantly, my interpretation of *consciousness* does not say how each of the dimensions of the conceptual space are to be scaled because this is to be determined according to our theoretical interests. In particular, I think that each of the approaches to *consciousness* that I have analyzed in the previous chapter determines a specific way to scale the conceptual space. For instance, an intentionalist could elaborate a scale for the intentionality dimension in terms of how abstract the contents of conscious states can be and suggest that the scales of the other two dimensions are relative to the first. Analogously, a phenomenalist could argue that the phenomenality dimension is to be scaled in terms of the intensity of the what-it's-like-ness of each conscious state and that the scales of the other two dimensions are relative to the first.

The aim of my interpretation of *consciousness* is to provide a common ground for the study of consciousness that unifies different approaches according to a single conceptual space, by specifying the conceptual level to which the study of consciousness belongs. For instance, we could interpret the intentional approach as a configuration of the conceptual space where the reflexive and the phenomenal characters of conscious states are determined by their presentational characters, which means that *reflexivity* and *phenomenality* are to be explained in terms of *intentionality*. Analogously, we could interpret the reflexive approach as a configuration of the conceptual space where *intentionality* and *phenomenality* are to be explained in terms of *reflexivity*. Finally, the phenomenal approach can be seen as a configuration of the conceptual space where *intentionality* and *reflexivity* are to be explained in terms of *phenomenality*. Notice that all three of these approaches assign a higher theoretical weight to one of the dimensions of the conceptual space, which does not mean that they eliminate the other two; what it means is that those other dimensions somehow depend on the fundamental one. I grant that each of these configurations of the conceptual space is a viable candidate for explaining aspects of consciousness, but I see no reason to privilege any of them as the best model to explain consciousness in general.

⁴⁸ It might turn out that intentionality, reflexivity and phenomenality are not properties of consciousness, but conceptualizations of our ways to identify it. However, if this were the case, that would not mean that my characterization of *consciousness* does not correspond to consciousness.

Since *consciousness* is constituted by *intentionality*, *reflexivity* and *phenomenality*, the dimensions of its conceptual space are irreducible. However, the irreducibility of the dimensions could be interpreted in two different ways⁴⁹. On the one hand, it could mean that all conscious states have all three features at the same time, i.e., that they are necessary and sufficient conditions for a state to be conscious. On the other hand, it could mean that every conscious state could be described in terms of the components of *consciousness* or combinations of them. Recall that I do not know what consciousness is, neither what its actual properties are, so I cannot assess the first option. In particular, since my view concerns our characterizations of *consciousness*, which are always determined by our chosen practices of conceptualization, the status of the features of consciousness or of the components of its concept as sufficient or necessary conditions varies according to those theoretical interests. Hence, the second reading is the one I am proposing, insofar as I am discussing the irreducibility of the dimensions of the conceptual space of *consciousness*, not of the properties of consciousness.

In my view, our attributions of consciousness are enabled by our ability to identify and conceptualize the characters exhibited by conscious states. Thus, when we describe a conscious state in terms of, say, *intentionality*, such description does not mean that the state cannot also exhibit reflexive and/or phenomenal characters⁵⁰. Consider that we identify the characters of another person's state indirectly. For instance, I cannot inspect the phenomenality of your current state, though I have learned some practices of conceptualization that would allow me to infer it, given the proper circumstances. In general, we cannot determine the specific characters of another system's state, unless we have the proper conceptual/theoretical tools to do so⁵¹. However, that does not prevent us from attributing consciousness. What I want to highlight by saying that the dimensions of the conceptual space of *consciousness* are irreducible is that a proper characterization of the concept cannot disregard any of its components.

To show the interdependence between the components of *consciousness*, let us consider the case of Mary (Jackson, 1986). When Mary thinks about the physiology of the

⁴⁹ I thank Luke Roelofs for calling my attention to this issue.

⁵⁰ I will elaborate on this thesis in the next chapter.

⁵¹ Since I have no training in neuroscience, for instance, I cannot attribute consciousness to a given system by analyzing what the screen of a brain scanner shows; by this I mean that I do not have the proper tools to draw the kind of conclusions that neuroscientists are capable of drawing by interpreting what the screen shows, not that they attribute consciousness solely on the basis of what the screen shows. However, other practices of conceptualization would allow me to attribute consciousness to a system by observing, say, its behavior. In fact, this is how we attribute consciousness to other humans in everyday life.

perception of the color of a ripe tomato in her black and white lab, she arguably entertains a conscious state that is different from the conscious state she entertains when she leaves the lab and sees a ripe tomato. The question is why they are different conscious states, considering that they are both about the color of ripe tomatoes. In my view, the answer is that the presentational and the phenomenal characters of her states are different⁵².

The states under discussion are: 1) Mary's thinking about the physiology of the perception of the color of a ripe tomato, and 2) her visual perception of a ripe tomato. In contrast to the way in which theorists normally address the problem of determining the difference between Mary's states⁵³, my interpretation of *consciousness* suggests the following. Regarding *intentionality*, we could say that 1) is about a piece of non-perceptual information concerning the physiology of the perception of the color of ripe tomatoes, while 2) is about the what-it's-like-ness of the visual perception of a ripe tomato. Hence, though they are both about the color of ripe tomatoes, their presentational characters are different. Regarding *phenomenality*, it is often argued that the states have a different what-it's-like-ness, considering that 2) is a perception of the color of a ripe tomato and 1) is not. Of course, this means that their phenomenal characters cannot be identical, not that 2) lacks phenomenality. Thus, we can distinguish both states in terms of *intentionality* and *phenomenality* by clarifying that their presentational and their phenomenal characters are different. Put briefly, Mary's states are different because they exhibit different characters, not because one has a character that the other lacks.

My interpretation of *consciousness* might lead to the thought that we are actually dealing with different concepts of consciousness and, consequently, with different subject-matters. However, I have argued that different characterizations of *consciousness* correspond to different practices of conceptualization, not to independent concepts of consciousness. As a matter of fact, one of the purposes of my interpretation of *consciousness* is to show that our understanding of the concept is more general than what most characterizations of it suggest. In particular, it seems that we can modify the way we understand *consciousness* to

⁵² Although both states could be described as conscious states in terms of *reflexivity*, it does not seem that we could then explain why they are different; after all, we can assume that Mary is aware of her own conscious states (Coleman, 2015a, p. 75).

⁵³ Many theorists, including Jackson, claim that the difference between Mary's states has to do with the phenomenology of her states, by apparently endorsing a phenomenal approach to *consciousness*. However, we have seen that *consciousness* is not identical to *phenomenal consciousness*.

fit different contexts and to highlight different features of conscious states which, after all, is what enables us to attribute consciousness.

My proposal also shows that *consciousness* is not identical nor reducible to any of its components. I think that the relation between *consciousness*, *intentionality*, *reflexivity* and *phenomenality* is similar to the relation between a spiderweb and the strings that preserve its shape and integrity: if I pull a single string, the rest of the web will be disturbed; if I rip one of the web's structural strings, the whole web will fall apart. Just as a proper characterization of any of the components of *consciousness* requires a proper characterization of the other components, a proper characterization of *consciousness* requires a proper characterization of all of its components and their relations⁵⁴.

A consequence of my interpretation of *consciousness* is that the difference between conscious states is not a matter of degrees or levels, but of ways to conceptualize the characters they exhibit⁵⁵. In particular, their differentiation has to do with alternative ways to configure the conceptual space according to different practices of conceptualization. Hence, I endorse the thesis that “the notion of a level of consciousness is ill-suited [...], for it implies that global states of consciousness can be ordered in terms of a single dimension” (Bayne et al., 2016, p. 412). Accordingly, Mary's perception of a ripe tomato, for instance, could be described as an intentional state or as a phenomenal one, depending on whether we are interested in the intentionality of her perception or in its phenomenality. In both cases, we would be dealing with different configurations of the conceptual space of *consciousness*, not with different concepts of consciousness.

My proposal also implies that *consciousness* is a system-based concept, i.e., a concept that needs to be customized according to the sensory/cognitive skills of each kind of system. Hence, a characterization of the kind of consciousness that a given system entertains can only be given by adjusting the conceptual space according to the system's sensory/cognitive skills. Put briefly, what such an investigation would require is a customization of the conceptual space of *consciousness*⁵⁶.

One possible objection against my interpretation of *consciousness* is that it seems epistemically vacuous, in the sense that my aim to provide a view that is metaphysically

⁵⁴ This is the circularity to which I have made reference above.

⁵⁵ I will elaborate on this thesis in the next chapter.

⁵⁶ In the fourth chapter, I will concentrate on this thesis to explore how my interpretation of *consciousness* could be employed to characterize non-human kinds of consciousness.

compatible with most current theories of consciousness might render my proposal insufficient for assessing those theories or for grounding an alternative explanation of our subject-matter. Relatedly, Crane has pointed out this risk in relation to the concept of mind: “we already have a rough conception of our subject-matter; what we are looking for is not an explicit definition, but a description of the mental phenomena which is sufficiently clear and detailed for us to recognize it as a description of the thing of which we have this conception” (Crane, 2001, p. 3). However, notice that I have not provided an “explicit definition” of consciousness, but an explication (Carnap, 1950, pp. 3–8) of *consciousness*. Furthermore, my interpretation only delimits the conceptual level to which the study of consciousness belongs, not its ultimate nature. It may be that we already have a “rough conception” of consciousness, but the constellation of competing theories available evinces that our subject-matter is not “sufficiently clear”; this is why I also believe that my proposal is compatible with different views regarding the ultimately correct analysis of *consciousness*.

Another possible objection could be that it is not clear whether my proposal implies that all conscious states exhibit all three characters or if one of them is enough for us to qualify a state as a conscious one⁵⁷. Indeed, to claim the first would mean that any state that does not exhibit all three characters is not a conscious state, while to claim the second would make my proposal extremely vague and impractical, for it would be very difficult to draw a clear line between conscious and unconscious states. However, my proposal does not entail either of the former theses, insofar as it is not about consciousness. My proposal concerns our ability to identify and conceptualize the characters exhibited by a given internal state⁵⁸, which is what enables the attribution of consciousness.

In the next section, I will elaborate on my interpretation of *consciousness* to show how its conceptual space can be adjusted to build what I call *models of consciousness*. As I will show, a model of consciousness is not a concept of consciousness, but a reconfiguration of the conceptual space that corresponds to a particular theoretical interest. To illustrate this,

⁵⁷ I thank Luke Roelofs and Sam Coleman for calling my attention to this issue.

⁵⁸ We might be able to identify the three characters in a given state, and we might only be able to identify a subset of them. What would change is our description of the state, not our attribution of consciousness. Consider that, although the intentionality of somatic experiences does not seem to be specifiable, we can legitimately describe them as conscious states in terms of, say, *phenomenality*. Analogously, that thinking about a mathematical equation does not seem to have a characteristic phenomenality does not imply that it cannot not be described as a conscious state. In the next chapter, I will explore different ways to describe conscious states that exhibit different combinations of the characters.

I will review some popular theories of consciousness and show that they can be seen as exemplifications of different models of consciousness.

2.2. Three models of consciousness

The interpretation of *consciousness* that I have introduced in the previous section resulted from my analysis of the features that are normally attributed to conscious states, namely, intentionality, reflexivity and phenomenality. However, my proposal says nothing about what consciousness is or what its actual properties are, so one might wonder how it could be implemented in a scientific context. Recall my claim that the approaches to *consciousness* that I have analyzed in the previous chapter are practices of conceptualization that correspond to different theoretical interests. Analogously, my interpretation of *consciousness* is a practice of conceptualization but, unlike those approaches, it does not rely on any metaphysical view regarding the nature of consciousness. Thus, what distinguishes my proposal from most current theories of consciousness is that it is intended to provide a conceptual background for the study of consciousness, not a specific metaphysical thesis about its nature, nor a definite characterization of the concept.

In this section, I will explore three ways to build models of consciousness by adjusting the conceptual space of *consciousness*. I will show that a model of consciousness is a reconfiguration of the conceptual space. Arguably, the following models could lead to novel methods to identify and measure consciousness, insofar as they can be adapted to different kinds of systems without the need to redefine *consciousness*. In particular, as I will show in the fourth chapter, the following models of consciousness could be employed to characterize kinds of non-human consciousness.

It is reasonable to think that different kinds of conscious systems entertain different kinds of conscious states, either because they have a different physiology or because they have different sensory/cognitive skills; to investigate this would lead to an explanation of the functional basis of the kind of consciousness that those kinds of systems entertain, i.e., of their awareness mechanisms. However, we should all agree that, if a system entertains conscious states, we should be able to characterize the kind of consciousness that it entertains in terms of a single concept, namely, *consciousness*. Hence, the study of consciousness must be relatively independent from the study of awareness, insofar as, say, pain must be painful regardless of whether it is experienced by a human, a dog, a Martian or an asteroid. I am not

saying that we can study consciousness without considering the mechanisms that enable it, what I am saying is that the study of consciousness belongs to a different conceptual level.

My suggestion is that the conceptual space of *consciousness* can be customized according to our theoretical interests and in consideration of the reasons we may have to attribute the paradigmatic features of conscious states to the internal states of a given system. Hence, such customizations of the conceptual space are determined by pragmatic criteria; similarly to what Dennett (1971) argues concerning intentional systems. For instance, if we have reasons to believe that a certain kind of system might entertain states that exhibit phenomenal characters but not reflexive ones, we might want to leave aside the reflexivity dimension to concentrate on the other two; this certainly would not imply that the system is unable to entertain states that exhibit reflexive characters.

There are, at least, three ways to adjust the conceptual space of *consciousness*, i.e., three models of consciousness. Each model involves different subsets of the dimensions of the conceptual space. I will call them *one-dimensional consciousness*, *two-dimensional consciousness*, and *three-dimensional consciousness*.

One-dimensional consciousness is a model based only on one of the components of *consciousness*. Consequently, there are three kinds of one-dimensional consciousness, one for each component. This model accounts for the conceptual background of those theories that privilege one of the discernible features of conscious states over the others. From the metaphysical point of view, this would be the model adopted by those theories that posit one of the features as the fundamental property of consciousness.

There are many theories of consciousness that could be interpreted as adopting a one-dimensional model of consciousness. Crane's (2001) *intentionalism* takes intentionality as the fundamental property of all mentality, which suggests that all conscious states exhibit a presentational character or at least that they could all be described in terms of *intentionality*. Thus, Crane's intentionalism can be interpreted as a thesis that explains reflexivity and phenomenality in terms of *intentionality*. Similarly, Block's (1995) concept of phenomenal consciousness implies that phenomenality is the fundamental property of consciousness, insofar as he takes *consciousness* to be equivalent to *phenomenal consciousness*. Thus, Block's thesis should be able to explain intentionality and reflexivity in terms of *phenomenality*. Finally, Coleman's (2015b) *quotational higher-order thought theory* could

be interpreted in terms of a one-dimensional model that privileges intentionality⁵⁹. According to Coleman, “[c]onscious states comprise sensory content plus subjective awareness of that content”, where “what is conscious –what is like something for the subject– is a *content*” (2015b, p. 2718). Thus, it would seem that the conceptual background of Coleman’s theory is a one-dimensional model that explains reflexivity and phenomenality in terms of *intentionality*. Note that these theories do not necessarily reduce or eliminate two of the dimensions of the conceptual space of *consciousness*. Adopting a one-dimensional model of consciousness is a strategy that concerns the conceptual background of a theory, not the ultimate nature of consciousness.

One-dimensional consciousness is a good model for explaining the characters exhibited by conscious states in isolation. Suppose, for instance, that we want to explain the intentionality of human belief. Endorsing a one-dimensional model would allow us to isolate and describe the contents of human belief, regardless of the other characters that human beliefs might exhibit. Accordingly, a theory like Crane’s would seem to be appropriate, insofar as it would allow us to isolate the intentionality of human beliefs from their other features. However, it should be noted that a one-dimensional model based on *intentionality* would be insufficient to explain, say, the intentionality of somatic experiences because, as I have argued in the first chapter, their intentionality seems to consist in their phenomenality. Thus, it would seem that we would need to endorse a different model of consciousness to account for conscious states that exhibit combinations of the characters.

Two-dimensional consciousness is a model that privileges two of the components of *consciousness*. Accordingly, there are three kinds of two-dimensional models, one for each pair of components: *intentionality* and *reflexivity*; *reflexivity* and *phenomenality*; and *intentionality* and *phenomenality*. In general, two-dimensional models of consciousness highlight the relationship between two of the dimensions of the conceptual space and allow us to describe some complex features of conscious states, insofar as they not only involve two dimensions, but also their relations.

In general, Higher-Order (HO) theories of consciousness could be interpreted as adopting two-dimensional models, insofar as they tend to “analyze the notion of a conscious mental state in terms of reflexive meta-mental self-awareness” (Van Gulick, 2018) which I interpret as the combination of *intentionality* and *reflexivity*. Despite their differences, most

⁵⁹ Although I grant that Coleman’s characterization of intentionality is not identical to mine.

HO theorists (Lycan, 1996; Rosenthal, 1993) claim that “what-it’s-likeness enters only when we become aware of that first-order state [the one that serves as the content of the higher-order state] and its qualitative properties by having an appropriate meta-state directed at it” (Van Gulick, 2018). Thus, HO theories seem to interpret phenomenality in terms of the other two discernible features, in the sense that they seem to suggest that phenomenality results from the system’s awareness of the phenomenal characters of its intentional states. Similarly, Kriegel’s (2009) *self-representational theory* could also be interpreted as adopting a two-dimensional model. According to Kriegel (2009, p. 2), “phenomenally conscious states have qualitative character [what I call *phenomenal character*] in virtue of representing environmental features [what I call *presentational character*] and subjective character [what I call *reflexive character*] in virtue of representing *themselves* [what I called *self-referentiality*]”. Thus, his theory should be able to explain phenomenality in terms of *intentionality* and *reflexivity*, insofar as the phenomenal characters of experience are characterized as “a matter of a complex compresence of qualitative character and subjective character” (Kriegel, 2009, p. 11)⁶⁰.

In any case, while theories that adopt a two-dimensional model seem to be able to account for more complex features of conscious states, they also require a more sophisticated theoretical framework; in particular, they need to provide an explanation of the relations between the two components involved. It should be noted that one-dimensional and two-dimensional models are adequate for their theoretical interests and that, for that reason, they should not be compared. What I am arguing is that the conceptual backgrounds of those theories can be described according to two different configurations of the conceptual space of *consciousness*, not that they are theories about different kinds of consciousness.

Finally, there is the three-dimensional model of consciousness, where conscious states are described in terms of all three components of *consciousness*. Of course, compared to the latter two models, three-dimensional consciousness is a highly complex model that requires considering the relations between the three dimensions of the conceptual space. Consequently, although this model might seem to be the most suitable for the study of, say, human consciousness, its implementation also corresponds to certain theoretical interests. Consider that we seem to be able to describe most of our conscious states according to one

⁶⁰ I thank Sam Coleman for helping me improve my interpretation of Kriegel’s theory.

⁶¹ Certainly, there are other theories that adopt two-dimensional models of consciousness. For instance, the view called *phenomenal intentionality* (Mendelovici, 2018; Pitt, 2004) seems to adopt a two-dimensional model based on *intentionality* and *phenomenality*.

or two of the components of *consciousness*, even if some of them somehow exhibit all three characters. For instance, when I analyzed Jackson's argument about Mary, I distinguished her thinking about the color of ripe tomatoes from her visual perception of a ripe tomato in terms of *intentionality* and *phenomenality*; although I acknowledge that both states also seem to exhibit a reflexive character, insofar as Mary is arguably aware of both states.

To my knowledge, there is no theory that could be interpreted as adopting a three-dimensional model of *consciousness*. However, I grant that most contemporary theories make reference to all three components of *consciousness* in different ways and according to different theoretical interests. Perhaps, the reason why no theory of consciousness seems to adopt the three-dimensional model is that it is the least attractive from the methodological point of view, insofar as it requires an account of aspects that might not be relevant for the study of a specific kind of conscious system. What seems to be the case is that our everyday attributions of consciousness are somehow based on this model, for we tend to describe conscious states according to different combinations of all three components of *consciousness*.

In contrast to most current theories of consciousness, my proposal does not rely on any metaphysical thesis about the nature of consciousness, but merely on a general characterization of *consciousness*. Hence, I am not claiming that intentionality, reflexivity and phenomenality are properties that are present in every conscious state, nor that they are the actual properties of consciousness. As far as I can tell, the nature of consciousness could be very different from my interpretation of *consciousness*, though they are meant to correspond.

In the next chapter, I will focus my attention on conscious states to present a way to identify and categorize them that entirely relies on my interpretation of *consciousness*. My aim there is not to provide an explanation of the nature of conscious states, but a conceptual framework that accounts for our different ways to describe and categorize them. In particular, I intend to show that conscious states can be classified into three broad categories: intentional, reflexive and phenomenal states. Accordingly, I will argue that those categories result from assigning a higher theoretical weight to each of the dimensions of the conceptual space of *consciousness*.

Chapter 3: Conscious States

An individual experience, taken by itself, is unanalyzable. Experiences, taken as a manifold, can be compared and ordered, and only through their order result the (quasi) constituents of the individual experiences.

Rudolf Carnap (2003, p. 149)

In the previous chapter, I have argued that *consciousness* can be represented as a three-dimensional conceptual space formed by *intentionality*, *reflexivity* and *phenomenality*. Furthermore, I presented three ways to reconfigure its conceptual space, which I called *models of consciousness*, and maintained that they can be employed to study consciousness according to different theoretical interests. However, since my intention was only to provide an interpretation of *consciousness*, I did not elaborate on our identification and description of conscious states.

In this chapter, I will introduce a view where conscious states are descriptions of internal states. The idea is that conscious states are those internal states that exhibit either of the paradigmatic characters⁶² to which *intentionality*, *reflexivity* and *phenomenality* point, or combinations of them. I will argue that conscious states can be categorized according to our ability to conceptualize the paradigmatic characters and their combinations. In particular, I will maintain that conscious states can be classified according to three general categories: intentional, reflexive, and phenomenal states. Given the conceptual nature of this work, those

⁶² Namely, the presentational, the reflexive and the phenomenal characters.

categories should be interpreted as labels that enable the description and classification of some internal states in terms of *consciousness*, not as an ultimate account of their nature.

Intentional, reflexive and phenomenal states are categories that represent three ways to describe a particular kind of internal states, called *conscious states*, i.e., they are three practices of conceptualization that emerge from my interpretation of *consciousness* and, as such, their implementation depends on pragmatic reasons. I will show that each of those categories assigns a higher theoretical value to a component of *consciousness*, while preserving their interdependence. Overall, what I have done so far is isolate the conceptual level that, in my view, corresponds to the study of consciousness, and the result could be thus illustrated:

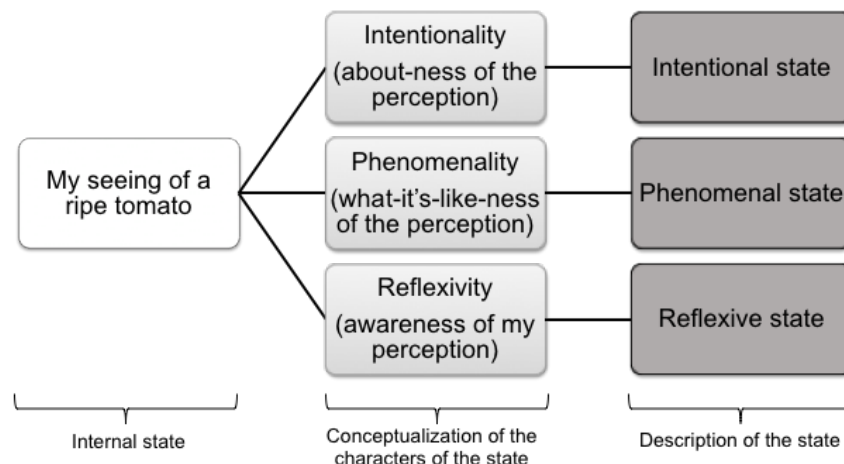


Figure 2: Ways to Describe Internal States

In the previous chapters, I focused on the light-colored areas. In what follows, I will elaborate on the dark-colored ones. Notice that the above illustration does not determine the nature of conscious states, neither that the categories cannot overlap. In fact, I will show that the same conscious state could be described according to different categories. Furthermore, conscious states could also be described as, say, functional, cognitive, or neural states, depending on our theoretical interests. My view implies that conscious states must be described according to *consciousness*, not that they are an independent kind of state; this is the foundation of my claim that the study of consciousness belongs to an irreducible conceptual level.

3.1. Intentional states

Many internal states exhibit a presentational character, which we identify as their about-ness, and their description as conscious states relies on how their about-ness is

conceptualized. Accordingly, intentional states are internal states that are mainly described in terms of *intentionality*. I will argue that the classification of some internal states as intentional states is a practice of conceptualization that consists in assigning a higher theoretical weight to the intentionality dimension of the conceptual space of *consciousness*, though that does not mean that intentional states only exhibit presentational characters. On the contrary, I will show that many of them also exhibit the other characters and that the conceptualization of their intentionality often depends on specifying the relation between their presentational characters and their other characters.

In my view, our ability to identify and conceptualize the presentational characters of a system's internal states enables the attribution of intentionality. I will not elaborate on how such identification is done because, presumably, each kind of system has its own mechanisms. However, at least from the human perspective, it seems that the identification of the presentational character of a state can be done directly, as when we identify the aboutness or the contents of our own states, or indirectly, by examining, for instance, the behavior and/or the brain activity of a system.

Once we have identified the presentational character of a given state, we can conceptualize it according to different theoretical strategies. It is very likely that newborns perceive whatever stimulates their senses, even if they are not yet able to represent what their perceptions target, and that is enough to conjecture that some of their internal states can be described as intentional states. Analogously, my niece's cat, Jack, is perfectly capable of distinguishing the food container from the water container, even if we have no reason to believe that it can represent those objects as "a food container" and as "a water container". So, in a sense, Jack seems to be able to entertain internal states that exhibit presentational characters, even if we have no reason to believe that it can represent the objects of those states. In contrast, since the internal states of, say, a digital thermometer do not seem to exhibit presentational characters, it would be odd to attribute intentionality to its states. Hence, the attribution of intentionality to some of the internal states of other systems somehow parallels the way in which we attribute it to our own states: we project our ways of identifying and conceptualizing the presentational characters of our own states into those of other systems.

In view that the dimensions of the conceptual space of *consciousness* are deeply interconnected, my suggestion is that intentional states are a category of conscious states that is grounded on a practice of conceptualization that privileges the intentionality

dimension and describes the other characters in terms of how they relate to the presentational character of a given state. According to that practice of conceptualization, intentional states could be distinguished by their objects or targets, i.e., by what they are about. Bear in mind that intentional states can be about objects, collections of objects, situations, events, concepts, propositions, beliefs, etc. Thus, my use of ‘object’ refers to the targets of intentional states⁶³, not to objects in the metaphysical sense of the term; after all, an event is not an object, though it can be what a state is about.

Many of our conscious states are about what I call *perceptible objects*. By ‘perceptible object’ I mean the presentation of whatever we internally or externally perceive. Perceptible objects, then, include the presentations of whatever stimulates our senses and those of the states of our bodies⁶⁴. Hence, whenever we perceive, we entertain states about perceptible objects. Notice that not all perceptible objects can be specified, as when we hear something that we cannot identify or when we have an internal sensation that we cannot describe. In consequence, that perceptible objects are the targets of perceptions does not mean that every perceptible object is necessarily represented in the mind. The intentional states that are about perceptible objects are what I call *perceptual states*.

Every system capable of entertaining states about what stimulates its senses or about at least a subset of the states of its body is an embodied conscious system, which means that only embodied systems can entertain perceptual states. Accordingly, the ability to entertain perceptual states is what we normally call *perception*. It is noteworthy that perception cannot be reduced to the sensory/cognitive mechanisms that underlie or enable it because, while it would seem reasonable to claim that a system that is able to adapt to its environment has sensory/cognitive mechanisms, that would not mean that the system perceives its environment in the sense of entertaining states that exhibit presentational characters. Self-driving cars, for instance, have sensors and mechanisms that enable their navigation, but that is not enough to conjecture that they entertain states about their surroundings in the relevant sense.

An appealing way to investigate perceptual states is to focus on states whose contents we can represent. For instance, a system that solves a maze could be described as a system capable of representing the maze, thus implying that the system entertains perceptual states.

⁶³ I thank Sam Coleman for helping me clarify my use of ‘object’.

⁶⁴ At least the states of our bodies of which we can become aware according to our mechanisms of bodily awareness.

However, this functionalization of perception in terms of sensory/cognitive skills would imply the reduction of the presentational characters of perceptual states to their basis, i.e., the reduction of consciousness to awareness⁶⁵. In contrast, my characterization of perceptual states includes states whose contents cannot be specified from a third-person perspective. In the first chapter, I argued that the intentionality of somatic experiences, moods, emotions, etc., consists in their phenomenality, which is a kind of presentational character that cannot be accounted for in terms of sensory/cognitive mechanisms. Consider that the intentionality of, say, a headache consists in its what-it's-like-ness, insofar as the target of the perception of the headache is its phenomenal character.

The presentational characters of some perceptual states can be described in terms of a relation between *intentionality* and *phenomenality*. Let us reconsider the system that solves a maze puzzle. It is not the fact that the system is able to find the exit what enables our attribution of perceptual states to it, but how it solves the puzzle. For instance, a system that learns to identify relevant cues to find the exit could be said to entertain perceptual states, insofar as the identification of those cues would suggest that the system has states about those cues and not mere awareness of its surroundings; in fact, mere awareness would not enable the system to identify cues. In a simplified manner, we could describe the system's internal states as perceptual states about its surroundings and the cues, which would be a description in terms of *intentionality*. However, it could be that the system's learning process involves the identification of phenomenal features, like odors and colors, as cues, in which case we would have to describe the intentionality of the system's perceptual states in terms of a relation between *intentionality* and *phenomenality*. Accordingly, we could describe the system's internal states as perceptual states that target phenomenal features, which would imply that those states also exhibit phenomenal characters. Thus, in my view, phenomenal features can be described as perceptible objects because they are what some perceptual states are about⁶⁶.

There are also perceptual states whose intentionality has to be described by means of *reflexivity*. In particular, those perceptual states that target the states of the system's body⁶⁷. In my view, what characterizes this kind of perceptual state is that their intentionality

⁶⁵ I will return to these methodological issues in the fourth chapter.

⁶⁶ That the features of the phenomenal characters of our experiences are often conceptualized as properties, namely, as what theorists call *qualia*, corresponds to a different practice of conceptualization that concerns the metaphysics of phenomenality, which exceeds the scope of this work.

⁶⁷ In the first chapter, we saw that the ability to entertain this kind of conscious state is called *bodily awareness*.

partially consists in their reflexivity, insofar as they involve the system's awareness of the states of its body. For instance, my perception of the orientation of my body involves my awareness of the relation between my body and my surroundings, which implies that I perceive my body in a way that fixes its relation to the rest of the perceptible objects around me. Thus, my perception of the orientation of my body can only be described by making explicit that its intentionality is at least partially constituted by my awareness of the state of my body and, consequently, its description relies on a relation between *intentionality* and *reflexivity*; possibly, the description would also involve *phenomenality*, under the assumption that experiences about our own bodies tend to exhibit phenomenal characters⁶⁸.

Not all intentional states are perceptual states, some are what I call *representational states*. This further kind of intentional state targets representations, which result from the exercise of our representational skills. Consequently, representations are not perceived, insofar as they do not directly stimulate a system's sensors. Since representations are themselves specifiable, they can be interpreted as the contents of representational states.

Some representational states are about what I call *individual representations*, which stand for perceptible objects or are formed by modifying other individual representations⁶⁹. Thus defined, individual representations include representations of perceptible objects, as well as representations of non-existent and fictional objects, like unicorns and Darth Vader. Allow me to exemplify what are representational states about individual representations.

I have never seen an Egyptian pyramid in person, though I have seen pictures and drawings of them. From those pictures and drawings, which are perceptible objects, I have formed a representation of what an Egyptian pyramid looks like. Notice that the pictures and drawings are perceptible objects that are conventionally⁷⁰ meant to represent the pyramid, but they are not the pyramid. In contrast, the representation I have formed from my perception of those pictures and drawings stands for the pyramid. Thus, whenever I imagine or think about an Egyptian pyramid, I entertain a representational state whose content is an individual representation of the pyramid. Indeed, a perceptual state about the actual pyramid and a representational state about it, are both intentional states about the actual pyramid, but

⁶⁸ I will address this in the next section.

⁶⁹ I call them *individual* precisely because they target single perceptible objects or composed representations, i.e., representations formed by modifying other individual representations. Of course, we can have several individual representations of, say, the same perceptible object, but my point is that they all share a common target.

⁷⁰ A system incapable of interpreting the purpose of a picture or a drawing would be unable to form a representation of what they depict.

they have different presentational characters, namely, the former is about a perceptible object and the latter is about a representation of it.

There are also representational states about non-existent and fictional objects, which are individual representations formed by modifying other representations. For instance, although I could form a representation of a unicorn by looking at, say, a drawing in a book, such representation does not stand for a perceptible object because the drawing depicts a representation formed by modifying other representations; arguably, one of a horse and one of a horn. Similarly, imagining a unicorn is to entertain an intentional state about an individual representation, not about a perceptible object. Regardless of how my representation of a unicorn is formed, it must be an individual representation because, even if I can modify it in several ways⁷¹, it always targets a unicorn.

The description of some representational states involves *reflexivity* and *phenomenality*. For instance, *intentionality* is not enough to account for some of the features of a representational state like imagining how it would be to have four eyes. Such representational state is not merely about a possible scenario, but about a possible scenario that concerns me, imagined from my current point of view, which is a feature that can only be accounted for in terms of *reflexivity*. Now try to imagine how it would be to face Darth Vader. If we limit our description of the state to the intentionality dimension, then we will not be able to explain the phenomenal character to which the state points because, after all, it is not simply a state about being in front of Darth Vader, but about the what-it's-like-ness of the encounter. Thus, imagining how it would be to face Darth Vader could be described as an intentional state about an individual representation of the what-it's-like-ness of facing Darth Vader, which would be a description of its intentionality in terms of *phenomenality*.

Some other representational states are about what I call *cultural objects*. In general, cultural objects are complex representations shared by a group of systems. Among this kind of objects, there are concepts, propositions, beliefs, opinions, etc., and together they give form to the way in which a system interprets its environment. Of course, a system's conceptualization of its environment does not have to involve every aspect of it. On the contrary, it seems that representational states about cultural objects target portions of the system's environment viewed from the perspective of the system and in accordance with the conventions that the system shares with the members of its group.

⁷¹ For instance, by representing the features of the unicorn differently.

Cultural objects can be formed by means of the conceptualization of perceptible and/or individual representations. What distinguishes cultural objects from individual representations is that, while the latter stand for individual objects, properties, events, etc., the former are labels or categories that group perceptible objects and/or individual representations according to the conventions of a group of systems.

Consider, for instance, the concept of cat, i.e., *cat*. While cats can be perceived or individually represented, *cat* cannot. I can have perceptual states and individually represent cats, but I cannot have an individual representation of *cat*. Certainly, whenever I perceive or imagine a cat, I am perceiving or imagining something to which *cat* applies, but I cannot have an individual representation of *cat* because it would be a representation of an individual cat, not of the concept. My point is that the characterization of *cat* involves a) the conceptualization of a kind of perceptible object and/or individual representation, and b) the adoption of a convention that concerns the application of the concept. This is why I believe that *cat* is a complex representation, i.e., a cultural object.

It might seem that representational states about cultural objects can be fully described in terms of *intentionality*. However, some of them can only be described by means of *reflexivity* and *phenomenality*. Consider, for instance, my belief that “cats make good pets”. Although my belief is about an opinion shared by many people, it also exhibits a reflexive character, insofar as it is about an opinion that I endorse. In other words, my belief has a reflexive character, inasmuch as it involves a propositional attitude that I hold towards the convention that “cats make good pets”. Hence, a full description of my belief must involve *reflexivity*, to properly capture its presentational character. Now consider my thought that “pizza is a delicious Italian meal”. While ‘pizza’ and ‘Italian meal’ point to cultural objects⁷², ‘delicious’ clearly points to the conceptualization of the phenomenal character of my experiences tasting pizzas. Hence, a proper description of the intentionality of my thought, apart from involving *reflexivity* in the above sense, must make reference to *phenomenality*.

Among cultural objects, some are what I call *theoretical objects*. These are a subset of cultural objects, insofar as they are also complex representations, but they play a key role in the construction of a worldview. In particular, theoretical objects are meant to transcend the particular worldview of a group of systems, in the sense that they are taken as universal

⁷² That these are cultural objects should be clear from the fact that, for instance, not everything that looks and tastes as a pizza is a pizza; at least not for an Italian person. Analogously, what qualifies as an Italian meal depends on conventional factors that are beyond the place where the meal is cooked or the person who cooks it.

concepts, propositions, etc. In other words, theoretical objects are assumed to be objective or universally true. Accordingly, the contents of theories are theoretical objects because they allegedly describe facts, irrespective of how a certain culture interprets those facts.

For instance, when we think about Newton's law of universal gravitation, we do it independently of our spatio-temporal coordinates and the conventions of the group to which we belong because the law is supposed to hold irrespectively of those factors. However, it is noteworthy that Newton's law is made true by his physical theory and by certain socio-cultural institutions, like the Royal Society⁷³. Hence, Newton's law of universal gravitation is a cultural object that becomes a theoretical object once his physical theory is legitimized by an authorized socio-cultural institution. Nevertheless, the law could also lose its status as a theoretical object, if the theory that makes it true gets discarded⁷⁴. Apart from physical theories, logical systems and mathematics are also constituted by theoretical objects; actually, their contents are supposed to be the most objective theoretical objects conceivable because of their apparent universality.

A representational state about a theoretical object differs from one about a cultural object in the sense that the former is meant to be universally true, while the latter could be, for instance, a belief based on the observation of a regularity that is ruled by a convention⁷⁵. Consider that the proposition "cats are animals" is backed by a biological theory, while the proposition "cats make good pets" is backed by cultural practices that might not be shared by all cultures or even by every member of a group⁷⁶. Therefore, what characterizes representational states about theoretical objects is that they constitute what we call *knowledge*.

The alleged universality of theoretical objects does not imply that representational states about them can be fully describe in terms of *intentionality* because some of them may also exhibit the other characters. In my view, since representational states about theoretical objects are a subset of representational states about cultural objects, they often have reflexivity and phenomenality. Consider the thought that "I am human". 'Being human' is a

⁷³ Although it might be argued that the Royal Society is a scientific institution and not a socio-cultural one, I take scientific institutions to be a kind of socio-cultural institution because they are conformed according to the worldview of a certain culture.

⁷⁴ A clear example of this change of status is the displacement of ether as an object within physical theories.

⁷⁵ Consider that astrology is grounded on the conventions of different groups and that, although it is supposed to be universal, it relies on cultural objects, like the relation between the position of certain celestial bodies and the personality of a person. Thus, when I think about my horoscope, I entertain a representational state about a cultural object.

⁷⁶ Although I grant that the universality and/or objectivity of theoretical objects is also relative to a group.

theoretical object supported by a biological theory, and it is meant to say something about the nature of some systems in our world; this accounts for part of the intentionality of my thought that “I am human”. However, the thought also exhibits a reflexive character that we can identify by realizing that it concerns my own nature. Thus, a description of the intentionality of the state must involve *reflexivity*; after all, the thought is not just about the nature of some systems, but about my nature. Furthermore, in this moment, while thinking about theoretical objects, I am entertaining an intentional state that has a distinctive, though probably unspecifiable, phenomenality. My current reflection on theoretical objects exhibits a phenomenal character that is part of its intentionality and, consequently, the intentionality of my current state would have to be described in terms of a relation between *intentionality* and *phenomenality*.

In this section, I have introduced a classification of conscious states that assigns a higher theoretical weight to the intentionality dimension of *consciousness*, i.e., I have described conscious states as intentional states. However, I have shown that the presentational characters of some intentional states can only be fully described by making reference to the other two components of *consciousness*, namely, *reflexivity* and *phenomenality*. My characterization of intentional states can be thus illustrated:

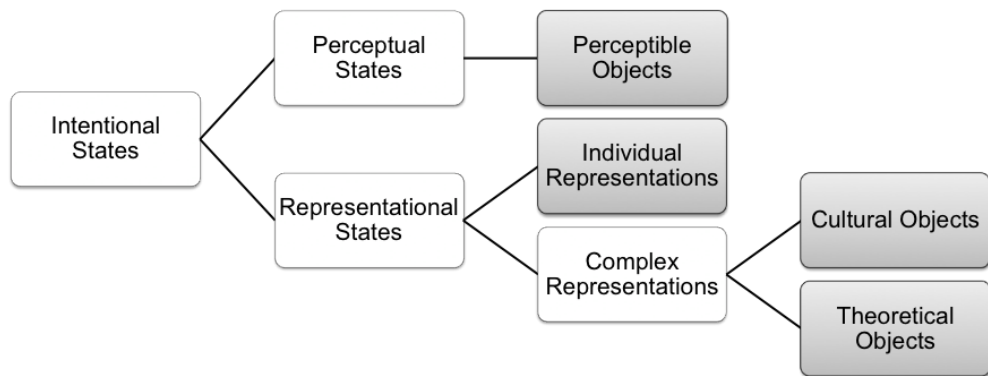


Figure 3: Intentional States

In the next section, I will show a different way to describe conscious states, namely, as reflexive states. In contrast to intentional states, which are descriptions of conscious states in terms of *intentionality*, reflexive states are descriptions of conscious states in terms of *reflexivity*, which means that reflexive states are not necessarily a different kind of conscious states, but a different way to describe them.

3.2. Reflexive states

Some internal states exhibit reflexive characters and their description often relies on how their reflexivity is conceptualized. Broadly speaking, the reflexivity of conscious states concerns the fact that conscious systems are often aware of themselves and their own mentality. Thus, those states that are mainly described in terms of *reflexivity* are what I call *reflexive states*. My suggestion is that reflexive states are a category of conscious states that is supported by a practice of conceptualization that assigns a higher theoretical weight to the reflexivity dimension of the conceptual space of *consciousness*, though they often exhibit the other two characters. I will show that the reflexivity of many conscious states can only be described by means of *intentionality* and/or *phenomenality*.

The reflexivity of conscious states has been characterized in several ways; for instance, as a *sense of mineness* (Guillot & Garcia-Carpintero, forthcoming) or *sense of ownership* (Peacocke, 2014). However, those characterizations already point to a relation between *reflexivity* and *phenomenality*. In my view, a system could be aware of its own mentality without experiencing any particular what-it's-like-ness, i.e., without having a “sense” of mineness or ownership, because reflexivity and phenomenality are conceptually distinguishable aspects of consciousness. Perhaps human reflexive states are always like something, but the connection between their reflexive and phenomenal characters is not necessary; arguably, reflexive characters may not even have a distinctive phenomenology (Bermúdez, 2011). Still, it is reasonable to believe that the relation between both characters produces what Coleman (2012, p. 145) calls a *phenomenal perspective*. In any case, these are ways to describe the reflexivity of at least some conscious states and it is noteworthy that, conceptually, they explicitly point to how some reflexive characters relate to some phenomenal characters.

I will now introduce a way to characterize the reflexivity of conscious states in terms of my interpretation of *consciousness*. Broadly speaking, I intend to argue that some conscious states can be categorized as reflexive states by analyzing how their reflexive characters manifest. In particular, I will hold that reflexive characters can be externalized or internalized. By ‘externalized’ I mean that the reflexive character is directed towards the outside of the system’s mentality, while by ‘internalized’ I mean that it is directed towards the system’s own mentality.

Note that the reflexive characters exhibited by states of bodily awareness⁷⁷ are directed outwards, i.e., externalized: I am aware of my own body, I can perceive and control it, and that awareness allows me to interact with my environment in different ways. Notice that the reflexivity of states of bodily awareness cannot be explained in terms of the mechanisms that enable it because a system could have, say, self-monitoring mechanisms, without being aware that what those mechanisms monitor are the states, or some of the states, of its own body. Many electronic devices have self-monitoring mechanisms, but they do not seem to be able to entertain states of bodily awareness, insofar as those mechanisms do not generate states that exhibit externalized reflexive characters. Furthermore, we are not aware of every state of our bodies, though our bodies have mechanisms that monitor many of their states: my body is said to be able to monitor my blood pressure, but I certainly cannot become aware of it. Hence, it would seem that the reflexivity of states of bodily awareness is externalized in the sense that it is projected towards the outside of the system's mentality⁷⁸.

States of bodily awareness often exhibit intentionality and phenomenality, which means that their description as reflexive states may involve *intentionality* and/or *phenomenality*, if our theoretical interests demand it. Consider the feeling of hunger. Although such state is more frequently described as a phenomenal state⁷⁹, it could also be described as a reflexive state, insofar as it concerns my awareness of something happening inside my body. To experience hunger requires the ability to be aware of the bodily state associated to it, which means that hunger is always felt from the perspective of the system⁸⁰. Accordingly, the feeling of hunger exhibits an externalized reflexive character. In consequence, the description of the feeling of hunger as a reflexive state would have to characterize the reflexivity of the state in terms of a relation between *reflexivity* and *phenomenality*, to account for the fact that the feeling is experienced from the perspective of the system. The reflexivity of the same state of bodily awareness, namely, the feeling of hunger, could also be described in terms of *reflexivity* and *intentionality*, insofar as the state is arguably about something that is happening inside the system's body. However, I have

⁷⁷ It should be obvious that only embodied conscious systems are capable of entertaining states of bodily awareness; after all, these are states characterized by the system's awareness of at least some of the states of its own body. Nevertheless, embodiment is not conceptually necessary for entertaining reflexive states: if there were conscious angels, they would certainly have no bodily awareness.

⁷⁸ This should not be taken as a metaphysical claim because it only concerns the conceptual distinction between a system's mentality and its body.

⁷⁹ I will address this kind of description in the next section.

⁸⁰ Of course, this applies to every sensation, emotion, feeling, etc., because they are all experienced from the point of view of the system.

argued before that the intentionality of this kind of conscious states consists in their what-it's-like-ness, so their intentionality would have to be described in terms of *phenomenality*. In any case, my goal is to show that the same state could be described in different ways, according to our theoretical interests.

In my view, a full description of a state like the feeling hunger would have to involve the three components of *consciousness*, given that it is possible to identify the three characters in it. However, depending on our theoretical interests, we could decide to assign a higher theoretical weight to a subset of the components or, to borrow Dennett's (1971) terminology, adopt a different "strategy". This adds to the general claim of this chapter, namely, that we can privilege a subset of the components to describe a conscious state, thus implementing different practices of conceptualization. Of course, my view does not imply that every state of bodily awareness exhibits all three character⁸¹. For instance, according to my analysis in the first chapter, somatic experiences do not seem to exhibit specifiable presentational characters because their intentionality, if any, points to their phenomenal characters. Thus, we could categorize them as reflexive states by describing them in terms of *reflexivity* and *phenomenality*. What seems to be the case is that every state of bodily awareness exhibits an externalized reflexive character and that they can be described as reflexive states.

Other conscious states exhibit internalized reflexive characters, i.e., a reflexivity that is projected towards the system's own mentality, and such reflexivity is often called *self-awareness* or *self-consciousness*⁸². Broadly speaking, the idea is that some conscious states involve the system's awareness of itself as something distinct; which is why I claim that they exhibit internalized reflexive characters. In a sense, the reflexivity of any state of self-awareness, or self-consciousness, could be described in terms of a relation between *reflexivity* and *intentionality*, insofar as they are all about the system's mentality. However, we should bear in mind that the intentionality of these states might not be representational, in the sense that I do not need to represent myself as a self⁸³ to entertain states about my mentality.

The reflexivity of states of self-consciousness could be described in terms of what Carruthers calls *the strong notion of self-consciousness*, insofar as they involve "higher-

⁸¹ More precisely, that we can identify all of them.

⁸² I take 'self-awareness' and 'self-consciousness' to be equivalent because they both point to the system's ability to entertain states about itself and its mentality.

⁸³ I could even deny the existence of selves, as independent entities, and still entertain states about myself.

order awareness of oneself as a self, as a being with mental states and a subjective inner life” (2003, p. 12). However, this characterization seems to imply that the system must have the ability to form a representation of itself as, for instance, a self or a subject. The concept of self is controversial and some theorists have claimed that it might not point to a distinct referent (Siderits et al., 2010). Still, even if *self* lacks a clear referent, a system could be able to entertain states of self-consciousness because their reflexivity seems to point to the system’s mentality, not necessarily to a representation of the subject to which that mentality is attributed. Thus, it might be argued that states of self-consciousness can be reflexive in two senses, some exhibit reflexive characters that consist in the system’s awareness of its mentality, some others imply the system’s ability to represent itself as the subject of that mentality. The former sense would involve the capacity to entertain states that exhibit internalized reflexive characters, while the latter would further require representational skills.

In my view, the reflexivity of states of self-awareness, or self-consciousness, consists in the system’s ability to entertain states about its own mentality, which means that the system may not be able to form a representation of itself as such-and-such to be self-conscious. Indeed, many human self-conscious states seem to require representational skills, but that does not mean that every human self-conscious state is also a representational state. The difference lies in how we describe their reflexivity, i.e., in the practice of conceptualization that we adopt. On the one hand, we could say that these are reflexive states because they exhibit internalized reflexive characters, which would be a description that mainly privileges the reflexivity dimension of the conceptual space of *consciousness*. On the other hand, we could describe them as reflexive states in terms of a relation between *reflexivity* and *intentionality*, by pointing out that they are states that target the system’s mentality.

It could also be argued that self-conscious states are like something for the system, i.e., that they exhibit phenomenal characters, in which case they could be described in terms of a relation between *reflexivity* and *phenomenality*. For instance, to think about what it is like to have a headache seems to be a reflexive state that concerns the phenomenality of my previous experiences of a headache. Self-conscious states could even be described in terms of the three components of *consciousness*. For instance, to think about how it felt to get married could be described as a reflexive state about the what-it’s-like-ness of my experience of getting married.

Some self-conscious states exhibit reflexive characters that are closely linked to their presentational characters. Particularly, theorists have argued that we are able to entertain states of self-knowledge, which are states that involve “knowledge of one’s own sensations, thoughts, beliefs, and other mental states” (Gertler, 2019). However, it is noteworthy that the reflexivity of states of self-knowledge is not different from that of states of self-consciousness, they just involve a different set of representational skills. In my view, states of self-knowledge are a subset of self-conscious states that result from the system’s adoption of propositional attitudes towards its own mentality. In particular, states of self-knowledge are reflexive states whose reflexivity is intertwined with their intentionality, insofar as they target the system’s mentality in specific ways. Thus, propositional attitudes are modes of being aware of our own mentality that are enabled by some of our representational skills and, in consequence, they must be described in terms of a relation between *reflexivity* and *intentionality*⁸⁴.

The reflexive characters of conscious states could certainly be manifest in other ways or, more precisely, they could be conceptualized differently. Here, I have suggested a characterization of the reflexive characters of conscious states that is based on the direction in which those characters point. In any case, to say that some conscious states are reflexive states is a way to describe some internal states in terms of *reflexivity*. In other words, my claim is that reflexive states are not an independent kind of conscious state, but a category that corresponds to a particular practice of conceptualization.

The latter classification of conscious states assigns a higher theoretical weight to the reflexivity dimension of *consciousness*, i.e., describes conscious states in terms of *reflexivity*. However, I have shown that the reflexivity of some conscious states can only be described by making reference to the other two components of *consciousness*. My characterization of reflexive states can be thus illustrated:

⁸⁴ Of course, states of self-knowledge might also have a what-it’s-like-ness, insofar as they are a subset of states of self-consciousness.

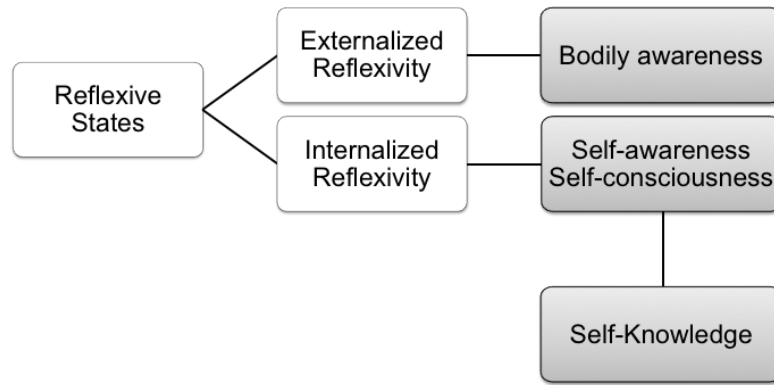


Figure 4: Reflexive States

In the next section, I will show a different way to describe conscious states, namely, as phenomenal states. In contrast to intentional and reflexive states, phenomenal states are descriptions of conscious states in terms of *phenomenality*, which means that they are not a different kind of conscious state, but a different way to describe them.

3.3. Phenomenal states

Many, if not all, conscious states exhibit phenomenal characters, though we may not always be able to identify or conceptualize them. In general, the phenomenal character of a conscious state is described in terms of what it is like to entertain that state, i.e., in terms of *phenomenality*⁸⁵. I call *phenomenal states* to those conscious states that are mainly described by conceptualizing their phenomenal characters. Although phenomenal states often exhibit either of the other two characters, their description broadly depends on how their phenomenal characters are conceptualized, thus making the description of the other characters dependent on the phenomenality dimension of the conceptual space of *consciousness*. In this section, I will argue that this way to categorize conscious states is a practice of conceptualization that assigns a higher theoretical weight to the phenomenality dimension of the conceptual space.

The conceptual strategy of describing conscious states as phenomenal states suggests that *consciousness* is equivalent to *phenomenal consciousness*. Many contemporary theorists (Brüntrup & Jaskolla, 2017; Coleman, 2015a; Goff, 2017) have argued that the properties of conscious states are phenomenal properties and that they are part of the fundamental reality. However, notice that my interpretation of *consciousness* does not necessarily entail that

⁸⁵ While many theorists characterize phenomenal characters as properties of conscious states, I only assume that they are discernible features of conscious states that can be identified and conceptualized according to *phenomenality*.

phenomenality is a property of consciousness because I have described it as a discernible feature, which means that it may not be independent from the other two features. Put briefly, my claim is merely that some internal states exhibit phenomenal characters. In any case, even if we grant that phenomenality is a property of consciousness, it seems that not all conscious states can be fully described in terms of *phenomenality* because they often also exhibit intentionality and/or reflexivity.

We seem to be able to identify the phenomenal characters of most of our conscious states and to conceptualize them in terms of what it is like to be in them. When we ask a child “where does it hurt?” or “how much does it hurt?” what we expect is enough information to be able to conceptualize the phenomenal character of the child’s experience. Similarly, when we ask, “how is the chicken?” at a restaurant, what we expect is a description of what it is like to taste it. Arguably, the phenomenality of conscious states is enabled by the system’s sensory/cognitive skills, but its description cannot be reduced to a description of those skills. That apples normally taste sweet to us does not mean that they must taste sweet to every possible conscious system; in fact, under certain circumstances⁸⁶, they may not even taste sweet to us. Hence, the phenomenality of conscious states must be described in terms of *consciousness*, not *awareness*.

While we are able to identify the phenomenal characters of many of our experiences, we cannot directly identify the phenomenal characters of the experiences of other conscious systems. I can describe the what-it’s-like-ness of tasting an apple by conceptualizing the phenomenal character of my experiences tasting apples, but I cannot describe what it is like for you to taste an apple by the same means. However, there are indirect ways to infer and describe the phenomenality of the experiences of another system. For instance, we could infer the what-it’s-like-ness of a system’s experience by associating its behavior to the phenomenality of the states that normally precede, underlie or follow such behavior⁸⁷. We could also analyze the neural states of a system while it presumably entertains a particular kind of conscious state and establish a relation between those neural states and the phenomenality normally attributed to the kind of conscious state under consideration⁸⁸.

⁸⁶ For instance, due to damage to the system’s taste receptors.

⁸⁷ Arguably, we employ this strategy when interacting with other humans and with those non-human systems to which we attribute at least some aspects of consciousness.

⁸⁸ This is surely an oversimplified description of the strategy that is regularly implemented in scientific contexts.

Since I take phenomenal states to be a category of conscious states that are mainly described in terms of *phenomenality*, I will present I way to describe them by analyzing how the conceptualization of their phenomenal characters relates to the other components of *consciousness*. In other words, I will show how conscious states can be described in terms of how *phenomenality* relates to *intentionality* and *reflexivity*.

From the conceptual point of view, there might be conscious states that can be fully described in terms of *phenomenality*. For instance, it seems that moods and emotions could be described in terms of what it is like to be in them, regardless of their causes and the specific sensory/cognitive skills of the system that experiences them. Consider the experience of surprise. The experience of surprise might target many different things, so it would be difficult to describe it purely in terms of *intentionality*⁸⁹. Furthermore, many kinds of systems, besides humans, seem to be able to experience surprise, so it is reasonable to assume that the feeling cannot be fully described in terms of *reflexivity*. While some features of the phenomenality of the experience of surprise might depend on the system's awareness of the experience or on its target, its phenomenality must always be relatively constant, insofar as every experience of surprise must be like something that can be described irrespective of the other characters. Thus, a description of the experience of surprise in terms of what it is like to feel surprised, i.e., in terms of *phenomenality*, would seem to be an adequate description of the relevant kind of state.

A description of a conscious state given only in terms of *phenomenality* would have to be a description in terms of what theorists call *phenomenal concepts*. I call these *pure phenomenal states*⁹⁰, not because they only exhibit phenomenal characters, but because their description relies entirely in the phenomenality dimension of *consciousness*. In my view, phenomenal concepts result from the abstraction of the what-it's-like-ness of certain paradigmatic conscious states, i.e., they are conceptualizations that disregard the other two characters. For instance, the concept of pain points to the characteristic what-it's-like-ness of certain experiences, irrespective of their intentionality and their reflexivity, because all experiences of pain seem to have a similar phenomenality. Thus, any conscious system capable of experiencing pain is a system capable of entertaining states that exhibit the distinctive phenomenal character to which *pain* refers.

⁸⁹ That we tend to associate a particular feeling of surprise to its target does not describe the phenomenality of the experience.

⁹⁰ It might turn out that there are no pure phenomenal states, but they certainly are conceptually possible.

Not all conscious states that exhibit phenomenal characters can be fully described by means of phenomenal concepts. In particular, there are conscious states whose description as phenomenal states requires making reference to *intentionality* and/or *reflexivity*.

Some conscious states exhibit presentational characters that modify or can be modified by their phenomenal characters, so their description must be given in terms of a relation between *phenomenality* and *intentionality*. I call these *intentional phenomenal states*. For instance, I have argued above that the categorization of somatic experiences as intentional states requires a description of their intentionality in terms of *phenomenality* because they are about what it is like to experience them. In contrast, now I am arguing that the same kind of experience can be categorized as a kind of phenomenal state, by inverting the practice of conceptualization employed to describe them⁹¹. Thus, somatic experiences can be considered phenomenal states by conceptualizing their phenomenality in terms of *intentionality*.

Intentional phenomenal states are conscious states that exhibit both phenomenal and presentational characters but are mainly described in terms of *phenomenality*. These are states that, apart from their identifiable phenomenal characters, have identifiable targets. For instance, what I called *perceptual states* above, namely, intentional states about perceptible objects, can also be described as intentional phenomenal states, by assigning a higher theoretical weight to the conceptualization of their phenomenal characters. Arguably, every perception has a what-it's-like-ness, so they could be categorized as phenomenal states that are about perceptible objects, in which case their descriptions would be focused on what it is like to perceive those objects.

Other conscious states exhibit reflexive characters that modify or can be modified by their phenomenal characters, so their description must be given in terms of a relation between *phenomenality* and *reflexivity*. I call them *reflexive phenomenal states*. In general, it could be argued that every experience is experienced by a system and that, consequently, they all have a certain reflexivity. However, their reflexive characters are not always identifiable or relevant for their description. Notice that the description of the phenomenality of pain is independent of how a particular system experiences it because every experience of pain must be painful, i.e., must exhibit the distinctive phenomenal character to which *pain* refers.

⁹¹ This is the reason why I have been claiming throughout this work that different approaches to *consciousness* correspond to different practices of conceptualization, not to different concepts of consciousness.

In a sense, reflexive phenomenal states are conscious states that have perspectivity, they are states about what it is like to entertain a certain conscious state from the point of view of a particular system or kind of system. In other words, these are conscious states whose phenomenal characters are modified by the system's awareness of them. For instance, while every experience of pain must be painful, the phenomenal character of an experience of pain can be experienced differently depending on its reflexive character. It often happens that football players get hurt without being aware of the injury and, by means of becoming aware of the injury, they experience the pain it causes in a different way. While it would be odd to claim that, before becoming aware of the injury, the player did not feel pain⁹², it seems that becoming aware of the injury modifies how the pain is experienced by the player. Thus, it would seem that the reflexive character of a given state can modify its phenomenal character. In the case of the football player, it is the awareness of the injury what modifies how the player feels the pain because the cause of the pain was already there and, presumably, the player was already in pain before becoming aware of the injury. In other words, the player's awareness of the injury changes the way in which the pain is felt, but it is not the cause of the pain. Thus, it seems that a proper description of the player's states, before and after becoming aware of the injury, would have to be given in terms of how the awareness of the injury modifies how the player experiences the pain, i.e., by means of a relation between *phenomenality* and *reflexivity*.

Similarly, it could be argued that the phenomenal and the reflexive characters of, say, self-conscious states are interrelated and that, consequently, self-conscious states could be described as reflexive phenomenal states⁹³. Self-conscious states clearly exhibit reflexive characters, which is why they tend to be described as reflexive states, but it could be argued that at least some of them have distinctive phenomenal characters. In this moment, I am thinking about the fact that I am the author of this work, and to entertain that thought is like something for me. Although the phenomenality of my current conscious state is difficult to describe, insofar as it involves many different experiences that I have had during the last two years, it should be clear that to think "I am the author of this work" is like something for me, precisely because of how my awareness of the thought modifies the what-it's-like-ness of entertaining such thought. Certainly, you can try to imagine what it is like for me to entertain the thought, but it is very unlikely that your state will exhibit a similar phenomenal character

⁹² Because the sensation of pain is not caused by the awareness of the injury.

⁹³ As I argued before, self-conscious states can also be described as intentional or reflexive states, by means of a relation between *intentionality* and *reflexivity*.

because it would have a different perspective, i.e., a different reflexive character, insofar as you are not the author of this work.

Finally, it seems that some phenomenal states can only be described in terms of the three components of *consciousness*. I call these *subjective states*, insofar as they are described as states about what it is like for a system to experience something. The difference between reflexive phenomenal states and subjective states is that the former may not exhibit identifiable presentational characters⁹⁴, while the latter do. In other words, subjective states are phenomenal states that exhibit identifiable presentational and reflexive characters. The idea is that the reflexive and the presentational characters of some conscious states can modify and be modified by their phenomenal characters.

Subjective states are of great interest to theorists because they seem to be the most complex kind of conscious state. However, notice that my approach implies that they are descriptions of some internal states given in terms of the three components of *consciousness*, not that they are a special kind of conscious state. In other words, my view does not entail that there are conscious states that are more complex than others. In a sense, all conscious states could be categorized as subjective states because, one could argue, conscious states must always be experienced by a system and presumably, be like something and about something. However, I believe that, depending on our theoretical interests, we may not need to describe every conscious state in terms of *intentionality*, *reflexivity* and *phenomenality*. For instance, a visual perception of a tomato, which is clearly a conscious state, could be described in terms of *intentionality* and *phenomenality*, even if the perception must be experienced by a system to be a conscious state, because the reflexive character of the perception may not be theoretically relevant⁹⁵. In other cases, a subset of the characters of a given conscious state may not be identifiable, which would prevent us from describing them as subjective states⁹⁶.

⁹⁴ Arguably, the player's experience of pain is not about the injury but about the sensation of pain. To say that the injury is what the sensation of pain is about would be to attribute the state an intentionality that derives from the conceptualization of its phenomenal character, i.e., to describe the state according to the practice of conceptualization that assigns a higher theoretical weight to *intentionality*, in which case the state would be described as an intentional state.

⁹⁵ From a scientific perspective, the description of the perception of a tomato does not need to make reference to the system's awareness of the perception because we can take for granted that there is a system perceiving the tomato; unless our interest is in investigating the sensory/cognitive mechanisms that enable the perception which, according to my approach, would be an investigation of awareness, not consciousness.

⁹⁶ Again, it seems that the presentational characters of, say, moods, emotions and somatic experiences cannot be identified, unless we adopt the practice of conceptualization that describes conscious states as intentional states, which amounts to a derivation of their intentionality by means of the conceptualization of their other characters in terms of *intentionality*.

Nevertheless, the visual perception of a tomato could be described as a subjective state under certain circumstances. Let us reconsider Mary’s perception of a tomato after leaving her black and white lab. Arguably, Mary is aware of her own mentality, so we could assume that she is capable of entertaining states that exhibit reflexive characters. Furthermore, since Mary is an unimpaired human, or at least this is what Jackson’s argument seems to suggest, we could conjecture that her sensory/cognitive skills are like ours and, consequently, that she must be able to entertain states that exhibit presentational and phenomenal characters. Hence, all else being equal, Mary’s perception of the tomato could be described as a subjective state, insofar as she must be capable of entertaining states that exhibit the three characters⁹⁷. Analogously, my present conscious state, sitting in from of my laptop thinking about subjective states while feeling a mild pain on my neck, seems to be a clear instance of a subjective state; in fact, I have just described it in terms of the three components of *consciousness*.

I have presented a way to describe conscious states that assigns a higher theoretical weight to the phenomenality dimension of *consciousness*. The idea is that some conscious states can be categorized as phenomenal states because they exhibit phenomenal characters that can be conceptualized either in terms of *phenomenality* or in terms of relations between *phenomenality* and the other components of *consciousness*. My characterization of phenomenal states can be thus illustrated:

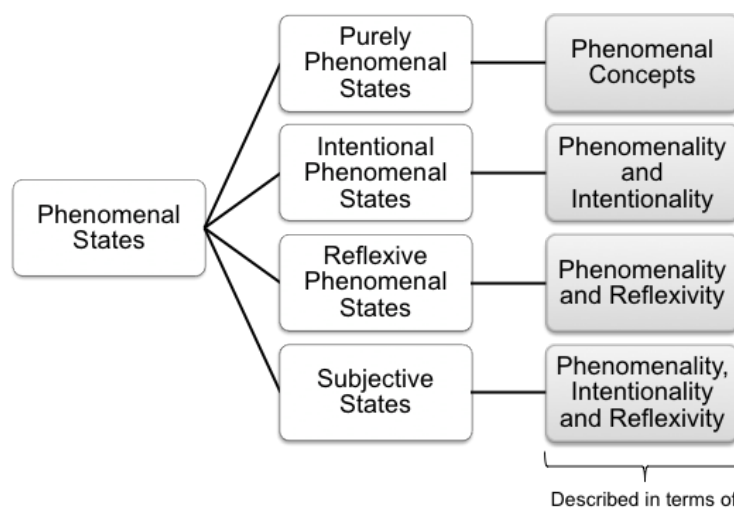


Figure 5: Phenomenal States

⁹⁷ Thus, in my view, the difference between Mary’s conscious states about the color of tomatoes in her black and white lab, and her visual perception of a tomato is that her visual perception exhibits a different phenomenal character, regardless of the fact that both states presumably target the same thing and that they arguably exhibit reflexivity.

In this chapter, I have argued that conscious states can be categorized according to three different practices of conceptualization, each of which assigns a higher theoretical weight to one of the dimensions of the conceptual space of *consciousness*. My goal was to show that conscious states are descriptions of those internal states to which we attribute the discernible features of conscious states and that their differences depend on how those features are conceptualized and categorized.

Chapter 4: Characterizing Non-Human Consciousness

Classical scientific definitions of intelligence use humans as a yardstick by which all other species are measured. According to these anthropocentric definitions, humans are always at the top of the intelligence rankings, followed by animals that look like us (chimpanzees, bonobos, etc.), followed again by other “higher” animals, and onward and downward in a league table.

Merlin Sheldrake (2020, p. 16)

In this chapter I intend to show how my interpretation of *consciousness* can be employed to analyze possible cases of non-human consciousness. Clearly, since my approach is mainly conceptual, I will not claim that there are non-human kinds of conscious systems. Instead, I will assess whether consciousness could be attributed to non-human kinds of systems by relying on my interpretation of *consciousness*. The difference is that, while the assertion of the existence of non-human conscious systems would have significant metaphysical implications regarding the extension of consciousness in reality, the attribution of consciousness is a practice of conceptualization that solely depends on how we apply *consciousness*⁹⁸.

In what follows, I will employ my interpretation of *consciousness* to assess whether consciousness could be attributed to dolphins and bees. Methodologically speaking, I will

⁹⁸ This is one of the reasons why I believe that my view is metaphysically compatible with most contemporary theories of consciousness, namely, because our attributions of consciousness are partially independent of our knowledge of its ultimate nature. Consider that, even if it turned out that consciousness does not exist, we would still attribute it to ourselves and perhaps to some non-human systems.

analyze scientific observations about how dolphins and bees behave under several ordinary and artificial circumstances, i.e., when facing tasks that they normally find in their natural environments and tasks that are completely novel to them. Of course, it might be argued that this methodology entails a sort of functionalization of consciousness⁹⁹, insofar as I will interpret their behavior in terms of *consciousness*. However, it should be noted that one thing is to claim that a certain kind of system is conscious because of how it behaves, and another is to interpret the system's behavior in terms of *consciousness*. While the first constitutes a metaphysical claim regarding the nature of the system, the second has the scope of providing an explanation of the system's behavior. Moreover, it is noteworthy that we regularly attribute consciousness to other humans by means of interpreting their behavior as evidence of their consciousness, i.e., by functionalizing consciousness. Hence, why should the same methodology be inadequate for attributing consciousness to non-human systems?

Before assessing the possibility of attributing consciousness to dolphins and bees, I would like to address a possible objection. In contrast to the approaches to *consciousness* that I have analyzed in the first chapter, my interpretation of the concept is not equivalent to *human consciousness*. Indeed, I have characterized *consciousness* in terms of our ability to identify and conceptualize three paradigmatic characters exhibited by human conscious states, but I do not think that our attributions of consciousness are determined by *human consciousness*. In any case, it could be argued that every attribution of consciousness is ultimately an attribution of human-like consciousness¹⁰⁰, which would mean that every possible kind of consciousness must be human-like. However, I do not believe that this is the case. That we attribute consciousness by applying a concept that is based on human consciousness only means that we interpret the world from our own perspective, not that the world is as we interpret it. In my view, what is human-like is the conceptualization of the features of conscious states, not the features themselves, which suggests that non-human conscious systems might entertain internal states that exhibit either of the characters in non-human ways.

The main purpose of this chapter is to introduce a way to study consciousness that does not necessitate a set of metaphysical claims and arguments to support it. The attribution of consciousness is a practice of conceptualization meant to describe and explain the behavior of a system or kind of systems by means of *consciousness*. In other words, the

⁹⁹ I thank Sam Coleman for calling my attention to this issue.

¹⁰⁰ I thank Luke Roelofs for calling my attention to this issue.

attribution of consciousness consists in the description of at least some of the internal states of a system in terms of *consciousness* as a means to account for the behavior of the system, or part of it. Certainly, not every behavior needs to be explained by attributing consciousness to the system. We have developed sophisticated methods and models that are capable of predicting how some systems behave or will behave under certain circumstances so, one might ask, when is the attribution of consciousness necessary? Why would such attribution lead to a better or more accurate explanation?

Perhaps someday we will have theoretical models powerful enough to predict even human behavior but, as many theorists have pointed out, such models may never be able to tell us why humans behave as they do; indeed, this seems to be the reason why we explain our behavior in terms of *consciousness*, instead of satisfying ourselves with mere causal accounts. Throughout this chapter I will build a case for the thesis that the attribution of consciousness to non-human systems might provide a more accurate explanation of their behavior. But what is a more accurate explanation? Accuracy is certainly a relative term, for it depends on our theoretical interests. Hence, I will argue that the attribution of at least some aspects of consciousness to dolphins and bees may lead to more accurate explanations, in the sense that such explanations seem to be simpler than the functional/computational models often employed by scientists. As I have been arguing, the attribution of consciousness is a pragmatical choice, a strategy that follows specific theoretical interests.

We normally explain some of our actions by attributing ourselves the ability to make decisions, instead of searching for an algorithm or theoretical model that mirrors those actions under similar circumstances. Sure, we can then employ theoretical models and algorithms to predict certain behaviors under similar conditions, but those models and algorithms do not explain, nor describe, why we act as we do. What explains our actions is precisely the attribution of the ability to make decisions. Analogously, I will argue that the attribution of consciousness to dolphins and bees provides a more accurate explanation of why they behave as they do. In consequence, I will not be arguing against the theoretical tools that scientists normally employ, for they may be more accurate to explain the functional/cognitive basis that underlies dolphin and bee behavior. Keep in mind that my main claim throughout this work has been that the study of consciousness belongs to an irreducible conceptual level, not to an independent metaphysical category. Thus, what follows is a speculation about why dolphins and bees behave as they do, by attributing them

the capacity to entertain at least some kinds of conscious states, not an argument in favor of the existence of non-human consciousness.

4.1. Dolphin¹⁰¹ consciousness

The scientific study of consciousness in non-human systems mirrors the study of human consciousness, especially regarding the conceptual background employed to identify the target of the study. For instance, theorists tend to differentiate the cognitive aspects of conscious states from what makes them conscious: “[i]t is essential to distinguish between “intelligence” (as problem solving) and “consciousness” (as wakeful alertness and conscious perception, including the perception of pain and pleasure)” (Baars, 2005, p. 12). I have expressed this distinction as one between *awareness* and *consciousness*. Broadly speaking, the idea is that a system could entertain cognitive states, i.e., have awareness, without being a conscious system, while it seems that every conscious system must be a cognitive system. Accordingly, theorists often claim that consciousness necessitates awareness, though awareness does not guarantee consciousness.

It seems that a system’s awareness mechanisms shape the way in which the system experiences, i.e., entertains conscious states, but the point is that consciousness cannot be fully explained in terms of its functional/cognitive basis. In my view, the reason why consciousness cannot be explained in terms of *awareness* is that *consciousness* and *awareness* are not the same concept and, more importantly, because the study of consciousness belongs to an irreducible conceptual level. Notice that I am talking about different conceptual levels, not different targets. Conscious states, understood as internal states to which we attribute the paradigmatic features of conscious states, could also be described as cognitive states, by means of *awareness*¹⁰². Although it could be argued that consciousness makes a cognitive difference¹⁰³, inasmuch as awareness and consciousness seem to be somehow intertwined, they must be studied from different conceptual levels.

In what follows, I will analyze some scientific observations regarding the way in which dolphins behave under natural and artificial circumstances, to assess whether we can attribute consciousness to them. The purpose of this analysis is to show how my interpretation of *consciousness* could be employed to explain dolphin behavior, not to assert

¹⁰¹ For simplicity, I will use the term ‘dolphin’ to refer to all members of the Delphinidae family.

¹⁰² Just as a conscious state could be categorized, for instance, as an intentional state or as a phenomenal state, depending on our theoretical interests and the practice of conceptualization that we adopt.

¹⁰³ I thank Sam Coleman for calling my attention to this issue.

that dolphins are, in fact, conscious systems. Broadly speaking, this section is inspired by the fact that

[s]cientific study has confirmed that bottlenose dolphins are large-brained, highly social mammals with an extended developmental period, flexible cognitive capacities, and powerful acoustic abilities including a sophisticated echolocation system. These findings have led some to ask if dolphins experience aspects of consciousness (Harley, 2013, p. 565).

At first glance, it seems that, despite the considerable distance between dolphins and primates in the phylogenetic tree, dolphins exhibit an impressive capacity for flexible behavior and adaptation, which is a sign that they have awareness mechanisms. Since the primary source of evidence for supporting this claim about dolphin awareness is behavioral analysis, I will now evaluate whether attributing them at least some aspects of consciousness would give us a better understanding of their behavior. In particular, I will argue that, while the attribution of awareness to dolphins is an empirical task, and there is a vast literature on the subject (Herman, 2006, 2012; Herzing & Johnson, 2015; Janik, 2013; Marino et al., 2007), the attribution of consciousness to them requires a conceptually different approach. As Harley points out,

[t]he existing evidence does not provide a convincing case for consciousness in dolphins. For productive scientific work on consciousness in dolphins (and other animals including humans), we need clearer characterizations of consciousness, better methods for studying it, and appropriate paradigms for interpreting outcomes (Harley, 2013, p. 565).

My interpretation of *consciousness* might aid the situation. In the second chapter, I have provided a definition of our subject-matter and argued that the conceptual space of *consciousness* can be modelled to account for any possible kind of consciousness. Hence, according to my characterization of conscious states in the third chapter, if we can describe some of the dolphin's internal states in terms of the components of *consciousness*, then we might have a case for attributing consciousness to dolphins.

4.1.1. The dolphin world

The first thing that we must take into account when we consider the possibility of dolphin consciousness is that consciousness, in general, implies perspectivity. Conscious systems are conscious of only a portion of their environment, which is determined by their awareness mechanisms. For instance, we know that the hearing range of dolphins is far

broader than the human (Roitblat, 2002, p. 183), but we also know that their sense of smell is considerably limited or even totally absent (Hanke & Erdsack, 2015, p. 56). In consequence, if dolphins are conscious, they are not conscious of exactly the same portion of the world of which we are¹⁰⁴. To properly assess the possibility of dolphin consciousness, we must describe the dolphin world in terms of their awareness mechanisms.

As far as we know, “dolphins obtain information about the identity, location, and characteristics of objects in their world by actively interrogating them using their unique biological sonar”; in particular, “[u]sing echolocation, dolphins can identify many characteristics of submerged objects, including size, structure, shape, and material composition” (Roitblat, 2002, p. 183). The dolphin biological sonar is a highly sophisticated sensory mechanism that cannot be reduced, nor compared, to our hearing. Moreover, to suitably describe the dolphin world we need to account for the integration between the dolphin’s hearing, sight and echolocation (Harley et al., 1996). Thus, although “[w]e cannot know what the dolphin’s subjective experience of perception is like, [...] we can know something about what it experiences by identifying the sensory dimensions that are available” (Roitblat, 2002, p. 186).

While there is no objective method to determine whether dolphins entertain phenomenal states, there are reliable ways to investigate if some of their internal states exhibit presentational and reflexive characters. Thus, I will adopt a two-dimensional model of consciousness, constituted by *intentionality* and *reflexivity*, to assess the possibility of attributing conscious states to dolphins. Let us start by considering the possibility of attributing intentional states to them.

We know that dolphins are capable of identifying and locating things in their aquatic environment, as well as distinguishing several characteristics of those things (Roitblat, 2002). However, the question is whether dolphins entertain internal states that exhibit presentational characters. Dolphins are not only capable of identifying objects by means of a single perceptual dimension, vision or echolocation¹⁰⁵, they are also capable of integrating information from both modalities (Harley et al., 1996; Herman et al., 1998; Pack & Herman, 1995) to make cross-modal identifications. Taken separately, each modality allows the dolphin to capture different features of its environment and to discriminate objects:

¹⁰⁴ Compared to ours, the dolphin world would have few or no smells, but many more sounds.

¹⁰⁵ The two main sensory dimensions of dolphins.

Using echolocation, dolphins can discriminate aspect-dependent objects (i.e., objects that produce different echoes depending upon the perspective from which they are echolocated) presented at different angles [...].

Dolphins are well known to use their biological sonar to obtain information about objects in their environment [...]. Dolphins can determine the material from which an object is made, the object's size, shape, and wall thickness [...].

Dolphins also have good vision in air and underwater [...]. Dolphins can visually discriminate objects that vary on several features at once. They can also distinguish 2-dimensional patterns that vary only in form (Harley et al., 1996, p. 165).

What these observations suggest is that dolphins can identify and distinguish what I called *perceptible objects*, independently of the sensory dimension with which they first perceived the object (Harley et al., 1996). For instance, they have the “ability to learn simultaneously about visual and acoustic properties of a set of stimuli” (Harley et al., 1996, p. 165), as well as the capacity to integrate visual and echoic information learned separately to improve their ability to discriminate objects. Furthermore, dolphins are capable of making cross-modal matchings to identify specific features of objects that are only available to one sensory dimension.

The fact that the dolphin could perform a cross-modal matching task suggests that her representational system may be hierarchically organized. For example, if the dolphin echolocates a sphere, she receives acoustic information about that sphere. If she is then required to use vision to choose that sphere from among a group of objects, she must have some way to connect her visual experience of the sphere with the acoustic information she had previously received in order to accomplish her task. A hierarchical system might allow her to tag the two experiences as being related because they both represent the sphere, or because they both represent a round object. This tag, whether it is based on simultaneous activation of the two representations or a connection with a third representation, is a form of hierarchical organization. That is, the identity of the object or object feature is on one level; the dolphin's different sensory experiences are on another level beneath it. Such a system might provide a mechanism through which the dolphin could create a stable representation of its world (Harley et al., 1996, p. 173).

According to these observations, the dolphin's ability to integrate multimodal information to identify and discriminate objects suggests that it “has an object-based representational system” (Harley et al., 1996, p. 164). Thus, by means of *intentionality*, we could conjecture that dolphins may be able to entertain perceptual states, i.e., internal states that exhibit presentational characters about perceptible objects. However, notice that my

description of perceptual states does not necessarily imply that the system represents perceptible objects, even less that it must be able to form “hierarchically organized” representations. The observations above suggest that dolphins perceive their environment and some objects within it, not that they represent those objects as such-and-such. For instance, we do not yet have reasons to believe that dolphins are able to represent isolated features of the objects they perceive, like roundness. It might well be that dolphins use some features of the objects they perceive as cues, but this does not mean that they represent those features in a specific way.

The ability to represent implies that the system must be able to recall, modify, relate and update its representations, which are complex cognitive skills. So now we must enquire whether dolphins have those skills. Theorists have observed that dolphins efficiently interact and adapt to changes in their environment, which suggests that they have cognitive skills, like attention (Ridgway et al., 2006), memory (Herman et al., 1989; Herman & Gordon, 1974), reasoning (Herman, 2006), and decision making (Johnson et al., 2015). Thus, it would seem that dolphins might be able to form representations after all. For instance, they seem to be able to remember, recall and update representations of objects that are not directly perceivable, or even absent (Herman & Forestell, 1985).

In response to a trainer producing the gesture for a named object followed by a “question” gesture, the dolphin pressed either a paddle to its right (to indicate “presence”) or a paddle to its left (to indicate “absence”). After learning the contingencies of the paddles using a small set of exemplars, the dolphin spontaneously responded accurately when novel queries were made about named objects (Pack, 2015, p. 185).

These observations evince that dolphins are not only aware of objects but may also be able to form representations of those objects to adapt to changes in their environment. Thus, it would seem that dolphins might entertain intentional states about non-perceptible kinds of objects, i.e., representational states about individual representations, considering that the capacity to determine the presence or absence of an object depends on the ability to form, store and recall a representation of an object that may not be directly accessible to the senses. In particular, dolphins exhibit the ability to associate representations to symbols, gestures or names, insofar as the production of a learned gesture by the trainer “evoked the mental representation –a “search image”– of the referenced item” (Pack, 2015, p. 185).

What about other kinds of representational states? Can dolphins form concepts? Can they have states about cultural or theoretical objects? Although it is more difficult to

investigate the possibility of entertaining these kinds of intentional states in non-human systems, some dolphin behaviors suggest that they might have the cognitive skills that underlie such complex representations. In particular, some of their behaviors might be explained by describing their internal states as intentional states about cultural objects.

Bottlenose dolphins at Laguna off the coast of Brazil have an unusual group-specific feeding technique which seems to date from 1847 and have been transmitted within a matrilineal community since at least three generations of dolphin are involved [...]. The 25–30 dolphins and local fishers follow a strict protocol –involving no training or commands from the fishermen– that allows the humans and dolphins to coordinate their actions. The dolphins drive fish into the nets of human fishermen, indicating as they do so by performing a distinctive rolling dive when the humans should cast their nets. The humans can also pick up from how much of the body comes out of the water on this roll an idea of how many fish are present –it is entirely unclear whether this cue is given intentionally or not– and then feed off the fish that are stunned or missed by the net (Rendell & Whitehead, 2001, p. 316).

Evidently, it all comes down to how we define culture. Compared to the richness of human culture, this kind of behavior might not be enough to say that there is anything like a dolphin culture. However, the fact that only a localized group of dolphins makes use of this feeding technique can be interpreted as a sign that a certain kind of knowledge is transmitted from one generation to the other. Thus, it would seem that describing the internal states of those dolphins as representational states about cultural objects would provide an explanation of their complex behavior. In particular, their internal states could be described as intentional states about a certain convention shared by the members of the group, which is precisely how I have defined cultural objects. In any case, my point is that a complex behavior like the one mentioned above does not seem to be explainable in terms that do not involve *consciousness*, nor in terms of basic cognitive skills, because, if that were the case, then we would observe similar behaviors in other groups of dolphins.

To have representational states about cultural and theoretical objects, a system must be capable of codifying information somehow and of forming concepts. The cognitive skill to form concepts, in turn, requires the ability to isolate relevant features or procedures and to be able to apply concepts under different and novel circumstances. One key feature of this cognitive ability is that the system must be able to identify semantic and syntactic differences, i.e., to learn and use a language.

Our studies of language comprehension have revealed capabilities in the dolphin for processing both semantic and syntactic information. [...] The dolphin is capable of understanding that changes in word order change meaning. It can respond appropriately, for instance, to such semantic contrasts as *surfboard person fetch* (take the person to the surfboard) and *person surfboard fetch* (take the surfboard to the person).

In these language studies, the dolphin demonstrated an implicit representation and understanding of the grammatical structure of the language (Herman, 2002a, pp. 278–279).

These observations evince that dolphins are capable of understanding semantic differences by identifying syntactic features, like the order of the elements of a sentence¹⁰⁶, which is a sign that they may be able to learn and transfer concepts according to changes in their environment. More interestingly, dolphins can recognize anomalous sentences and extract only the semantically and syntactically correct sequences from longer sentences without additional training (Herman et al., 1993). Furthermore, they are also highly efficient in providing judgements of numerosity (Yaman et al., 2012).

After a dolphin was trained with limited exemplars to respond to whichever of two arrays contained the greater number of items, it transferred the “more” concept to novel numerical contrasts. The reward contingencies were then reversed, requiring the dolphin to respond to the array containing the least number of items. After training with one numerical contrast, the dolphin readily applied the new “choose less” rule at performance levels of 85 percent or greater to other contrasts, including some it had responded to previously by choosing “more” (Pack, 2015, p. 183).

The complexity and plasticity of these dolphin behaviors, I suggest, is best explained by means of *consciousness* because the fact that they can invert a concept without any special training evinces that they are not just responding to stimuli. Admittedly, that dolphins seem to be capable of learning concepts and use them in novel circumstances is not a direct proof that they can have representational states about theoretical objects. However, it must also be said that the abilities to learn and apply concepts, like numerosity, are often at the base of the capacity to entertain intentional states whose contents are those concepts. For this reason, I suggest that attributing dolphins the ability to entertain intentional states about concepts might enable a better understanding of the complexity of their behavior.

¹⁰⁶ They are usually sequences of gestures, not words.

Although an assessment of the possibility to attribute intentional states about cultural and theoretical objects to dolphins would require further research, the complexity of their behaviors suggests that they might be able to entertain some intentional states, or at least that some of their internal states could be described in terms of *intentionality*. In particular, the attribution of intentional states to dolphins would allow us to claim that there is a dolphin world, i.e., that dolphins interact with their environment by means of perceiving and representing objects around them. In other words, some dolphin behaviors evince that they have a perspective and that they can adapt to their environment by means of perceiving and representing it according to their sensory/cognitive skills. My intention is not to assert that dolphins definitely entertain intentional states, but to suggest that an explanation of some of their behaviors would require the description of their corresponding internal states as states that exhibit presentational characters.

Having a perspective does not only depend on the ability to entertain intentional states. To have a perspective, a system must also be able to identify itself within its environment, as well as its own mentality, which means that the system must be able to entertain states that exhibit reflexive characters.

4.1.2. *The social dolphin*

Dolphins are highly social systems that live in what theorists call *fission-fusion societies*, “where group structure (i.e., both size and composition) may change on timescales of hours to minutes” (Würsig & Pearson, 2015, p. 89). The complexity of this kind of social organization suggests that dolphins are able to develop and keep track of a variety of social relations and bonds. Before examining dolphin social behaviors and relations, we must assess whether they have the cognitive skills necessary to enable reflexive states.

Dolphins have an outstanding ability to use their bodies to solve complex tasks. Their capacity to navigate in an omnidirectional environment and to plan trajectories and actions is also noteworthy. Furthermore, dolphins are capable of identifying and deliberately using specific parts of their bodies (Herman et al., 2001). In fact, researchers have investigated

[...] the dolphin’s awareness of its body parts by associating different human gestures with nine different body parts and then testing the dolphin’s comprehension of different symbolic sequences that requested the dolphin to use these body parts in different ways, including ways that were novel [...]. Overall, the dolphin performed at near-ceiling levels and was able to immediately transfer its use of named body parts successfully to novel objects and to objects positioned in novel ways. The dolphin’s

consistently high levels of performance demonstrated that it understood the body part symbols as representing its own body parts (Pack, 2015, p. 189).

These observations suggest that dolphins are aware of their own bodies and that they can move and use their body parts deliberately. The fact that dolphins can learn to identify and use their body parts could be explained by means of attributing them the ability to entertain states that exhibit at least the distinctive reflexive character of states of bodily awareness. In particular, the ability to associate gestures to their own body parts to perform specific actions does not seem to be explainable in terms of elemental cognitive skills. Consider, for instance, that the gestures are fully symbolic, in the sense that they do not point towards the body parts they name, so the dolphin must learn to interpret them correctly, which implies that dolphins have complex cognitive skills.

Moreover, since the above observations are about the dolphin's ability to learn to name and use parts of its body and considering that they appear to entertain at least some intentional states, it seems that an accurate explanation of the abovementioned behavior could be given in terms of a relation between *reflexivity* and *intentionality*. To perform the action indicated by the human trainer, the dolphin must be able to interpret the gestures, make the correct association to its body parts and determine how to act which, in my view, should be explained by attributing the dolphin the ability to entertain conscious states about parts of its own body, insofar as they exhibit the reflexivity of bodily awareness and the intentionality of perceptual states.

Impressively, dolphins are also capable of learning self-referential symbols and interpret them to solve difficult problems. Beyond their ability to identify and use their body parts to follow commands, dolphins are also very efficient at imitating and mirroring the behavior of others.

[F]or example, if the human raises his/her leg in the air, the dolphin will raise its tail, and if the human moves his/her head side to side or up and down, the dolphin will do the same. To accomplish these imitative acts, the dolphin must relate its body image to the body plan of the demonstrator, creating analogies, if necessary, for those behaviors demonstrated by a human (e.g., the dolphin's tail is analogous to the human's leg) (Herman et al., 2001, pp. 251–252).

Although the dolphin's abilities to imitate the actions of other dolphins are well documented (Herman, 2002b), it is clear that its capacity to mirror human actions involves far more complex cognitive skills. The fact that dolphins can interpret the movement of a

human leg as a tail movement, or the movement of an arm as the movement of a pectoral fin, seems to suggest that they are aware of themselves as individuals, insofar as they must be able to interpret the movements of the human in terms of their own bodies and according to their aquatic environment. For instance, when the human leans back to raise a leg, the dolphin turns on its back and raises its tail outside the water, which evinces their ability to interpret the human's action to determine an analogous movement. In my view, such behavior enables the attribution of reflexive states whose reflexive character is externalized, in the sense that the dolphin interprets its perception of the human action outside of itself, thus reproducing the action in a way that mirrors the way in which the dolphin perceives the movement from its perspective.

Dolphins are excellent imitators and part of their natural forms of association depend on their ability to imitate and create different patterns of behavior. For instance, a dolphin can “reliably repeat or not repeat its previous behavior”, which indicates that it can maintain “a mental representation of the behavior last performed” (Herman, 2002a, p. 280). The ability to deliberately repeat a specific behavior suggests that dolphins can form and update representations of themselves and their actions to determine their next actions. Moreover, they can also decide to perform a completely novel action (Herman, 2002b). These observations not only suggest that dolphins can entertain representational states about individual representations, namely, representations of themselves, but also reflexive states, insofar as those representations result from their awareness of themselves. Thus, it seems that we can describe at least some of their internal states as conscious states that result from relations between presentational and reflexive characters, which would be descriptions in terms of *intentionality* and *reflexivity*.

What about reflexive states whose reflexive character is internalized? Can dolphins be aware of their own mentality? One of the most popular methods aimed at investigating a system's self-awareness is the so-called *mirror self-recognition* (MSR) experiment, and Reiss & Marino (2001) have shown, though inconclusively, that dolphins might have the sensory/cognitive skills to track visual changes in their reflection on the mirror, as well as environmental changes. However, the method has clear limitations regarding the assessment of dolphin self-awareness because “in the dolphin we have a species whose natural environment is markedly different from that of the terrestrial animals usually studied in self-recognition experiments”, especially because “[i]t is difficult to see how, in the wild,

dolphins would be likely to encounter and learn to use reflecting surfaces as a source of information about self or others” (Loveland, 1995, p. 255).

At this point, we can turn back to the issue of dolphin culture. A complex social behavior, like the one exhibited by a group of dolphins off the coast of Brazil, entails that each member of the group has to learn and follow a specific set of conventions and rules. A sophisticated feeding technique, like the one exhibited by that specific group of dolphins, is not a merely mechanical procedure, but a series of coordinated actions learned and passed from one generation to the other.

Culture seems to be the kind of thing that requires both self-awareness and social awareness, for each individual must be able to pinpoint its role within the group and act according to the conventions and rules of the group. Hence, the dolphin’s abilities to imitate others, to produce novel actions, and its complex social behaviors could be explained by describing their internal states as states of self-awareness, or self-consciousness, in the sense that those behaviors seem to require the ability to entertain reflexive states about its own mentality. Notice that to determine its role within the group, a dolphin must be able to identify itself and to represent the relations it holds with other members of its group, which involves complex cognitive skills that are supposed to enable states of self-awareness. Thus, their complex social behaviors could be explained by attributing them the ability to entertain conscious states that exhibit internalized reflexive characters.

Evidently, the attribution of reflexive states to dolphins is highly speculative, insofar as there is no direct method to investigate their ability to entertain internal states that exhibit reflexive characters¹⁰⁷. However, my point is not that there is conclusive evidence that dolphins are able to entertain reflexive states, but that some of their behaviors could be explained more accurately by describing their internal states in terms of *reflexivity*. Thus, what I have argued so far is that there might be good theoretical reasons to attribute at least some intentional and reflexive states to dolphins, in light of the complexity of their behaviors and the evidence of their complex sensory/cognitive skills.

4.1.3. Can we attribute consciousness to dolphins?

There seems to be no doubt that dolphins have awareness, insofar as they exhibit behaviors that evince complex cognitive skills. However, the question is whether at least

¹⁰⁷ Although it must be acknowledged that there is also no direct method to investigate the reflexivity of human conscious states. Fortunately, language often serves to communicate the reflexivity of our conscious states, but this method is not available to dolphins.

some of their internal states can be described in terms of *consciousness*. As I have argued throughout this work, the identification of complex sensory/cognitive skills is not enough to attribute consciousness because conscious states further exhibit characters that can only be described in terms of *consciousness*. Thus, the attribution of conscious states to dolphins would have to prove that their behaviors are better explained by means of *consciousness*, which would motivate the hypothesis that dolphins might be conscious systems.

So far, I have argued that some dolphin behaviors can be explained by describing some of their internal states according to a two-dimensional model of consciousness based on *intentionality* and *reflexivity*. In other words, I have shown that some dolphin behaviors could be explained by attributing them intentional and reflexive states. Moreover, I have argued that some of their internal states could be described in terms of relations between *intentionality* and *reflexivity*. Nevertheless, I have argued that the intentionality and/or the reflexivity of some conscious states can only be described by means of *phenomenality*, so I will now investigate whether some dolphin internal states exhibit phenomenal characters. Although there is no direct method to identify the phenomenal character of a conscious state, there are indirect ways to do so. For instance, it would seem that a system capable of entertaining perceptual states would also be able to entertain states that exhibit phenomenal characters. After all, it is reasonable to assume that every perception has a what-it's-like-ness.

In general, we have no trouble in attributing phenomenal states to ourselves and other humans. Nevertheless, when it comes to non-human systems we hesitate, probably because we mostly rely on three human-based ways to infer the phenomenal character of a state, namely, neuroimaging, behavioral analysis and verbal reports¹⁰⁸. On the one hand, we should discard verbal reports, insofar as we cannot communicate with non-human systems in a way that is precise enough to take a verbal report as evidence of the phenomenality of a system's state¹⁰⁹. On the other hand, although neuroimaging seems to be a more reliable method, an analysis of the neural states of a system does not provide definite proof of its consciousness,

¹⁰⁸ My claim that these are human-based methods is motivated by the fact that they are mainly employed to infer features of human conscious states.

¹⁰⁹ Arguably, a parrot that pronounces the phrase "I am in pain" is not reporting the phenomenology of its current internal state. Analogously, dolphins cannot say "let's take five, I'm bored of these experiments", neither cross their fins in a sign of discomfort.

especially, of the phenomenology of its states¹¹⁰, because we do not know yet how consciousness relates to its apparent neural basis. What about behavioral analysis?

Although we often take some patterns of behavior as signs of the phenomenal characters of our conscious states, it must be acknowledged that the connection between those characters and their behavioral manifestations is conceptually contingent. In fact, we do not always behave in the same way when we entertain conscious states that are conventionally associated to certain patterns of behavior. Nevertheless, behavioral analysis could still be an indirect method to attribute phenomenal states to non-human systems, by following the same practices of conceptualization that we apply when analyzing our own behavior. Notice that the latter does not imply the functionalization of consciousness in terms of behavior because we are dealing with attributions of consciousness that have specific theoretical interests, not with assertions of consciousness. In other words, my proposal concerns the description of some internal states in terms of *consciousness* to explain behavior, which is not the same as taking certain patterns of behavior as evidence of consciousness.

I have shown that dolphins are very good at identifying and discriminating objects based on different features, like shape, size, orientation, etc., to argue that their internal states could be described in terms of *intentionality*. Thus, it seems that dolphins can form representations of individual objects, which is evinced by their ability to track the presence or absence of those objects in their environment. Dolphins also seem to be able to form complex representations of their environment. Consider that, without a complex representation of the environment as a frame of reference, dolphin's would not be capable of developing a feeding technique like the one observed off the coast of Brazil, or to use marine sponges on the rostrum as a protective tool for looking for prey in the substrate (Würsig & Pearson, 2015, pp. 91–92).

According to my view, the sensory/cognitive skills observed in dolphins, as well as the way in which they interact with their environment, suggests that they might be able to entertain states that exhibit phenomenal characters. In particular, the discrimination of features like shape and size suggests that dolphins are able to identify phenomenal features

¹¹⁰ According to what I have argued throughout this work, the study of consciousness and the study of its neural basis belong to different, irreducible, conceptual levels. Thus, neuroimaging can only serve to infer the phenomenology of conscious states once a direct relation between the paradigmatic features of conscious states and those of neural states has been conclusively established. Unfortunately, theorists do not universally agree on how to do this.

of their perceptions to modify their behavior. Thus, it would seem that at least some of their behaviors could be explained by describing their corresponding internal states in terms of *phenomenality*. Of course, this does not prove that dolphins can identify and conceptualize phenomenal features, like roundness, but that is not a requirement for entertaining states that exhibit phenomenal characters. My point is that dolphins exhibit sensory/cognitive skills that are best explained by attributing them the ability to entertain states that have phenomenality. Otherwise, how could they intentionally learn to discriminate the shape of an object, if their internal states did not have a what-it's-like-ness? Hence, to explain the complexity of some dolphin behaviors, it seems that the corresponding internal states should be described in terms of at least relations between *intentionality* and *phenomenality*.

Similarly, an explanation of the highly social behavior of dolphins might also imply a description of their internal states in terms of *phenomenality*. For decades, theorists have hypothesized that each dolphin has what they call a *signature whistle*, which is supposed to “broadcast their identity and location to other members of their social group” (Lammers & Oswald, 2015, p. 118). Although the precise role of these whistles is still unknown, “dolphins have been observed to mimic the signature whistles of other dolphins within a group [...], leading to the hypothesis that these whistles may be used to establish and maintain contact between individuals, particularly mother-calf pairs” (Lammers & Oswald, 2015, p. 118).

The way in which dolphins use their signature whistles evinces complex cognitive skills and it would seem that an explanation of those uses could be given by describing their internal states as reflexive states, insofar as the dolphin has to be aware of itself to be able to communicate its identity and location to other members of its group. Likewise, to properly interpret the function of a whistle, the other members of the group would have to be aware of their own identity and location, as well as having a representation of the dolphin emitting the whistle. Notice that the latter description of the internal states of dolphins engaging in this kind of interaction is given in terms of a relation between *intentionality* and *reflexivity*. Now, does this further suggest that dolphins might entertain phenomenal states concerning the bonds they have with the other members of their group? After all, they are able to recognize specific whistles and to adjust their behavior accordingly. Briefly put, could dolphins entertain emotions?

Although there is no definite evidence that dolphins entertain emotions, it seems that some of their social behaviors could be explained by describing their corresponding internal

states as phenomenal states; specifically, as what theorists call *affective states*. For instance, “[m]others will discipline their calves using tactile behavior such as making contact to the flank of a misbehaving calf with their rostrum, and even pinning their calf to the floor in more serious cases” (Kuczaj & Winship, 2015, p. 206), while “[p]ectoral fin rubbing appears to play a role in reconciliation between individuals after aggressive interactions and may also serve to reduce conflicts in juvenile-adult female relationships” (Kuczaj & Winship, 2015, p. 207).

Notice that the observations above already attribute affective states to dolphins, insofar as they imply that certain behaviors can be taken as evidence of their ability to entertain internal states that exhibit phenomenal characters. In contrast, I have argued that the connection between a certain pattern of behavior and consciousness is conceptually contingent. Hence, what we need is to show that a description of the internal states of dolphins in terms of *phenomenality* would enable a more accurate explanation of their behaviors.

From the observations mentioned above, we can conclude that dolphin social behaviors involve complex cognitive skills, insofar as they suggest that dolphins might be able to form bonds with specific members of their group and that they modify their behavior according to those bonds. Thus, it might be that those behaviors can be explained by means of a description of their internal states in cognitive terms, i.e., as resulting from their social skills. However, the question is whether the complexity of their social behaviors can be explained in terms that do not involve *consciousness*.

In my view, the attribution of phenomenal states to dolphins could explain the way in which they form and update representations of the members of their group. For instance, it would seem that a reconciliatory behavior could be explained by describing the corresponding internal states as exhibiting phenomenal characters that are associated to the representation that a dolphin has of itself, the other dolphin, and the bond between them. Such a description would explain, among other things, why dolphins engage in reconciliatory behaviors with specific members of their group and not with any other system. Notice that I am suggesting an explanation of the complexity of dolphin social behaviors in terms of *consciousness*, not a necessary connection between a certain pattern of behavior and consciousness.

According to the observations that I have reviewed in this section, there are good theoretical reasons to attribute consciousness to dolphins. By describing some dolphin

behaviors in terms of my interpretation of *consciousness*, I have shown that the complexity of those behaviors could be explained by attributing them the ability to entertain states that exhibit at least presentational and reflexive characters. Furthermore, by means of a two-dimensional model of consciousness based on *intentionality* and *reflexivity*, it seems that at least some dolphin internal states could be described in terms of *phenomenality*, insofar as the intentionality and the reflexivity exhibited by some of their internal states suggests that they may also entertain states that exhibit phenomenality. Of course, this does not definitely prove that dolphins are conscious systems, but my aim was to show how my approach could lead to a more accurate explanation of their behavior.

Briefly, it seems that some dolphin internal states could be described in terms of *intentionality* and *reflexivity*, as well as relations between them, which enables the attribution of consciousness to dolphins. Furthermore, since I have argued that the components of *consciousness* are interrelated, there might be reasons to conjecture that dolphins can also entertain internal states that exhibit phenomenal characters. Thus, by means of a two-dimensional model of consciousness, we could describe the phenomenal characters of at least some dolphin conscious states by conceptualizing them in terms of how they relate to the intentionality and the reflexivity of those states. In my view, the attribution of consciousness depends on how we conceptualized the characters exhibited by some of the internal states of a system, which means that my proposal concerns a practice of conceptualization based on my interpretation of *consciousness*, not the assertion that the system is a conscious system in the metaphysical sense.

4.2. Bee¹¹¹ consciousness

Traditionally, bees have been seen as simple systems whose behavior can be fully explained in functional/computational terms, or at least this is what traditional ethology suggests. Their small brains and their limited repertoire of behaviors is often held as evidence that they are natural *automata* which, in turn, suggests that there is no reason to attribute them any aspect of consciousness¹¹². However, theorists have recently argued that some bee behaviors could be explained by attributing them at least some aspects of consciousness. In this section, I will review some observations that might motivate a different approach to their

¹¹¹ Most of the observations that I will review are about honeybees and bumblebees. For simplicity, I will use the term ‘bee’ to refer to both, though I grant that there might be differences that would deserve a distinction.

¹¹² Since this reasoning is often generalized to all kinds of insects, my approach could also be adapted to study the possibility of attributing consciousness to other kinds of insects.

behavior. Before analyzing those observations, allow me to discuss one of the main reasons why theorists tend to reject the possibility of explaining insect behavior, in general, by attributing them at least some aspects of consciousness. According to the so-called *Morgan's Canon*:

In no case is an animal activity to be interpreted in terms of higher psychological processes, if it can be fairly interpreted in terms of processes which stand lower in the scale of psychological evolution and development (Morgan, 1903, p. 59).

My main concern about Morgan's Canon is that there is no obvious way to interpret it. On the one hand, it is not clear what a higher psychological process is, insofar as the same internal state could be described according to different practices of conceptualization. Consider that conscious states could also be described as psychological, cognitive, informational, functional or neural states, depending on our theoretical interests. On the other hand, what qualifies as a "fair" interpretation of a system's behavior? The fairness of an interpretation is always relative to a certain frame of reference. For instance, if human consciousness is our frame of reference, then the observation of bee behavior will hardly lead to the hypothesis that they might entertain conscious states because, compared to us, bees seem to have a significantly limited repertoire. Finally, Morgan's Canon implies that we must always choose the simpler interpretation available but it often happens that we do not know "which explanation is truly the simpler" (Perry & Barron, 2013, p. 19155).

My goal is to evaluate whether the description of at least some bee internal states in terms of *consciousness* would enable a more accurate explanation of their behavior. In contrast to Morgan's Canon, I will argue that the attribution of some aspects of consciousness to bees could lead to a simpler explanation of their behavior, without attributing them higher psychological processes. To assess the possibility of attributing consciousness to bees, I will employ a one-dimensional model based on *intentionality*. The reason why I have decided to employ said model is that the apparent simplicity of bee behavior suggests that most of their internal states result from how they interact with their environment, instead of the implementation of highly complex sensory/cognitive skills. Hence, if there is evidence that suggests that bees are aware of their environment, then we could address the hypothesis that some of their internal states could be described as intentional states.

4.2.1. *Beyond traditional ethology*

According to traditional ethology, bees are systems that, compared to any mammal, have a reasonably small brain and exhibit a limited repertoire of behaviors. However, they live in complex environments that require, among other things, the ability to avoid predators and to forage efficiently under constantly changing circumstances. In what follows, I will evaluate whether they exhibit the sensory/cognitive skills necessary to support the hypothesis that at least some of their internal states could exhibit presentational characters. In other words, I will investigate if bees are aware of their environment to determine whether at least some of their internal states could be described in terms of *intentionality*. Thus, the whole point of considering the possibility of bee consciousness is to challenge the claim that their behavior can be fully explained in functional/computational terms. Recent studies suggest that bees might not be mere *automata*. For instance,

[t]o forage successfully, a bee has to learn and remember the color, shape, and fragrance of flowers that are bountiful in these nutrients [nectar and pollen], and also how to get to them. Because the flowers that are in bloom are likely to change every few days, the bee needs, and has evolved, an impressive ability to learn (and re-learn) colors, odors, shapes, and routes quickly and accurately (Srinivasan, 2010, p. 268).

There is no doubt that bees are highly efficient foragers, but that does not necessarily entail that they are aware of their environment. For instance, from the point of view of traditional ethology, their foraging behaviors could have resulted from a set of evolutionary adaptations, which would mean that their identification and memorization of colors, shapes, etc., could be hardwired processes. To make the further claim that bees perceive their environment, we must investigate whether they can modify their behavior according to atypical circumstances because this would imply that they have legitimate sensory/cognitive skills, i.e., awareness mechanisms.

One way to investigate whether bees can deliberately modify their behavior according to changes in their environment is to test their capacity to learn from sensory stimuli, insofar as “[f]lexibility of behavior based on learning may indicate some cognitive capacities on the part of the animal” (Giurfa & Lehrer, 2005). Accordingly, if we can find reasons to believe that bees can learn from sensory stimuli, then it might be the case that they have the sensory/cognitive skills that are often said to enable or underlie internal states that exhibit presentational characters.

Recent research suggests that bees might have the cognitive skills that are at the base of what I called *perceptual states*. For instance, there are reasons to believe that they discriminate colors and odors to navigate. In a series of experiments (Giurfa et al., 1994, 1996, 1997), bees were able to learn that a certain color or odor at the entrance of a Y-maze indicated the direction to a reward. More interestingly, they were also able to invert the association between stimuli and direction in transfer experiments. These observations indicate that bees can detect and learn to use features of their environment to solve atypical problems. Moreover, there is evidence that bees can make cross-modal associations, like linking an odor to a color, or shapes to tactile information (Solvi et al., 2020). Perhaps these findings are not too impressive, given that the survival of bees partly depends on their ability to identify viable sources of food by discriminating different sensory cues. However, these observations show that bees can learn to make atypical associations.

Clearly, natural food sources are compounds of stimuli offering many cues simultaneously. Although bees may learn different cues in different tasks, they always use the most salient one, guiding them most reliably to the goal [...]. This is valid not only for cues belonging to the same sensory modality (e.g. vision), but also for cues related to sensory modalities as different as odour and colour (Giurfa et al., 1997, p. 241).

The fact that bees can learn to identify specific features to solve complex problems, as well as to invert associations when needed, suggests that they may have the sensory/cognitive skills necessary to modify their behavior in a non-mechanical manner, which could be taken as a sign of awareness. In particular, if bees were mere *automata*, they would hardly be able to make novel associations or cross-modal associations of novel stimuli. Thus, an explanation of their behavior seems to require at least the attribution of awareness, which enables the hypothesis that they might also be able to entertain perceptual states. Of course, this does not necessarily mean that bees can form representations of colors, odors, shapes, etc., and even less that they “perceive color [...] as a distinct sensation” (Srinivasan, 2010, p. 269), inasmuch as they could be capable of distinguishing colors without having sensations.

A different approach might give us a better insight into how bees employ their apparent sensory/cognitive skills. At first glance, we must acknowledge that, “[i]n their daily lives, bees are required to remember a number of different patterns and their properties, such as the shape of the nest or hive, shapes representing nectar-bearing flowers, and shapes of important landmarks on the way to the food source and back” (Srinivasan, 2010, p. 272).

Thus, if it turns out that their foraging and navigational behaviors can be explained by means of *intentionality*, we might have a case to conjecture that those behaviors are best explained by describing their internal states in terms of *consciousness*.

There is no doubt that bees are able to identify colors, shapes, odors and patterns (Giger & Srinivasan, 1995), so it seems reasonable to speculate that they might be able to perceive their environment. In other words, to explain the way in which bees interact with their environment, we could describe at least some of their internal states as perceptual states. Although we must grant that their perceptual states cannot be like ours because their sensory/cognitive skills are adapted to their environment, which greatly differs from ours.

Since bees are able to solve complex problems and to navigate their environment by identifying specific features, they could also be able to form representations of those features. However, it would seem that, if bees are able to form representations, those representations cannot be based on the identification of objects or features of objects, like most of our representations. In contrast to the way in which we represent our environment, bees could be said to represent features of their environment as directional cues (Horridge, 2005). Given the significant differences between their environment and ours, we should expect that, if bees have representational skills, those skills must be adapted to the way in which they perceive their environment. As Wystrach & Graham (2012) point out,

[...] the ability to recognize an object independently of its visual surroundings is necessary only if that object can be displaced and needs to be recognized in various locations, which by definition makes the object unreliable for navigation. Therefore we can wonder why animals would accept the processing cost of isolating and recognizing individual objects for navigational purposes, when simple egocentric views can encompass the overall structure and layout of a scene and underpin economical and robust navigation (Wystrach & Graham, 2012, p. 17).

Accordingly, if we attribute the ability to form representations to bees, we must take into account that, since their perceptual states have to be described according to their sensory/cognitive skills, their representational states are to be described in terms of how they appear to identify what they perceive. In particular, we should describe their perceptual states as intentional states about the landscape they perceive, as a unit, no about what I called *perceptible objects*¹¹³.

¹¹³ In a sense, we may be inclined to say that the intentional states attributed to bees could be about what Pylyshyn (2001) calls *proto-objects*, but the use of this term would be inappropriate, insofar as it takes human perception as its frame of reference.

One reason to believe that bees can form representations of their environment comes from the fact that they are very reliable navigators (Menzel & Giurfa, 2001). In particular, bees travel considerable distances to forage without losing their orientation, even in the dark (Chittka et al., 1999). For instance, recent scientific observations suggest that forager bees may be able to build some kind of representation of the surroundings of their nest or hive before traveling longer distances.

Upon first leaving the nest or hive, a new forager performs a ‘learning walk’ or ‘learning flight’, where a carefully orchestrated series of loops and turns allows her to learn the visual surroundings from perspectives that will be useful on subsequent return journeys [...]. Then, when the forager finally leaves the vicinity of the nest to forage, she is safely connected to it because of her path integration (PI) system. With PI, odometric and compass information are continuously combined such that at all times during a journey the forager has the approximate direction and distance information required to take a direct path home [...]. By remembering the coordinates of a successful foraging patch, the forager can also use PI to chart food-bound routes [...] or pass information to nestmates (Wystrach & Graham, 2012, p. 14).

These observations provide good reasons to hypothesize that bees can build, store and update representations of their environment to improve their foraging activities. More interestingly, it seems that forager bees have individual preferences and routes (Woodgate et al., 2016), which is a behavior that may be explained in terms of *consciousness*. Thus, a more suitable explanation of the behavior of forager bees could be based on the description of their internal states as representational states, insofar as such an interpretation would enable a more accurate explanation of their navigational preferences. Evidently, these observations do not prove that bees entertain representational states, but they seem to imply that the view of traditional ethology is insufficient.

According to the observations reviewed so far, we do not seem to have reasons to believe that bees could entertain intentional states about complex representations, though it remains a matter of empirical research. For instance, it would be interesting to investigate whether insects exhibit referential behaviors. In line with this hypothesis, Zhang & Srinivasan (1994) observed that, after training, bees can detect objects that are camouflaged or virtually invisible. However, that is not enough to consider that bees can form complex representations; at best, we could claim that they can learn to detect cues with which they were already acquainted under challenging circumstances. To further claim that they use

those cues referentially, for instance, we would have to investigate whether they can determine which cues are present and which are not.

The ability to entertain representational states about complex representations seems to be associated to the capacity to form concepts. In a broad sense, concepts allow conscious systems to solve complex problems and to rapidly adapt their behavior to novel situations. So, the ability to form concepts implies that the system must have complex cognitive skills like decision-making, context learning and problem solving. While there are no reasons to believe that bees understand concepts, they seem to be able to learn concept-like rules. In particular, it appears that they can learn distinctions like ‘same-different’ (Giurfa et al., 2001), ‘above-below’ (Avarguès-Weber et al., 2011), and ‘left-right’ (Avarguès-Weber et al., 2012); although it is still plausible that bees learn to make those distinctions without understanding concepts. However, theorists have argued that the ability to learn to follow a concept-like rule may be functionally equivalent to the ability to form concepts.

Learning of the concepts [...] was demonstrated through the protocols of delayed matching to sample (DMTS) and delayed non-matching to sample (DNMTS), respectively [...]. Honey bees foraging in a Y-maze [...] were trained in a DMTS experiment in which they were presented with a changing non-rewarded sample (i.e., one of two differently colored disks –‘color group’– or one of two different black-and-white gratings, vertical or horizontal –‘pattern group’) at the entrance of a maze [...]. The bees were rewarded only if they chose the stimulus identical to the sample once within the maze. Bees trained with colors and presented in transfer tests with black-and-white gratings that they had not experienced before solved the problem and chose the grating identical to the sample at the entrance of the maze. Similarly, bees trained with the gratings and tested with colors in transfer tests also solved the problem and chose the novel color corresponding to that of the sample grating at the maze entrance (Giurfa, 2013, p. 289).

While it might be that bees are not simply responding to stimuli in a mechanical manner, this does not justify the further claim that they have the cognitive skills to form concepts. That bees seem to be able to understand the task does not imply that they understand the concept involved. Perhaps a more plausible interpretation would be to say that bees learned to solve what we describe as concept-like problems by identifying relevant cues. Thus, in my view, the way in which bees behave under these atypical circumstances is not enough to claim that they can form concepts, because concepts are complex representations that may not be needed to solve matching tasks.

In a similar line of reasoning, we could assess whether bees can count, inasmuch as it requires the cognitive skill to numerically manipulate representations, which in turn would suggest that the system might be able to entertain representational states about complex representations. Recently, Skorupski and colleagues (2018, p. 7) argued that “[b]ees and most probably some other insects show a basic numerical competence, which may be limited to around four items”. However, they also acknowledged that the evidence is inconclusive because bees could be able to complete numerical tasks by non-numerical means¹¹⁴. In consequence, it seems that there is no reason to support the hypothesis that bees might be able to entertain representational states about complex representations, like concepts and numbers.

So far, I have shown that at least some bee behaviors could be explained by describing their corresponding internal states in terms of *intentionality*. In particular, I have argued that some of their internal states could be categorized as perceptual states and maybe even as representational states about individual representations; though the targets of those states have to be identified according to their sensory/cognitive skills, which differs from how I have identified them in the previous chapter. However, according to my interpretation of *consciousness*, *intentionality* is only one of the components of the concept, so I shall now investigate whether some bee internal states could also be described in terms of the other components of *consciousness*. Of course, since I am implementing a one-dimensional model based on *intentionality*, if some bee internal states exhibit any of the other two characters, namely, the reflexive and/or the phenomenal, I must describe them in terms of their relation to the presentational character of those states. Let us start by evaluating whether some bee internal states might exhibit reflexive characters.

Insects, as well as many other non-human systems, seem to have bodily awareness. In other words, insects seem to be able to entertain internal states that exhibit what I called an *externalized reflexive character*. For instance, bees can learn to use their bodies to pull a string attached to an artificial flower that would otherwise be inaccessible (Alem et al., 2016), which suggests that they may be aware of their bodies and that they can deliberately use their body parts to solve complex problems.

String pulling is a challenge that bees rarely find in their natural environment, so the fact that they can learn to solve the problem can be interpreted by means of attributing them

¹¹⁴ Similarly, it has been argued that ants are able to measure distance without actually counting steps (Wittlinger et al., 2006).

states of bodily awareness. Consider that, to solve the problem, bees need to be able to associate the string pulling to the movement of the flower towards the edge of the board that covers it, which requires the ability to orient their bodies in the appropriate direction and to determine how to grab and pull the string. However, we should acknowledge that “solving a string pulling task spontaneously is a relatively rare occurrence in bumblebees and might either reflect an unusually explorative “personality” in these individuals or simple “luck” in the process of random exploration” (Alem et al., 2016, p. 5).

It might seem that we should not interpret the string-pulling behavior as evidence of the ability to entertain reflexive states. Nevertheless, since we are dealing with a complex behavior that could be described in terms of *intentionality*, insofar as it targets the acquisition of the nectar by unnatural means, the presentational characters of the corresponding internal states could be described by attributing bodily awareness to the system, which is a conceptualization of the state’s apparent reflexive character. Briefly put, the string-pulling behavior could be explained by describing the bee’s internal state in terms of a relation between *intentionality* and *reflexivity*.

Although bodily awareness is the conceptualization of the reflexive character exhibited by many internal states of a variety of systems, including bees, the fact that some observations suggest that bees are bodily aware does not imply that they can entertain reflexive states. However, my implementation of a one-dimensional model based on *intentionality* suggests that string-pulling could be described in terms of a relation between *intentionality* and *reflexivity*, insofar as the behavior seems to exhibit identifiable presentational characters that can only be described by attributing a sort of reflexivity to the system’s state. Consequently, the description of some bee internal states in terms of relations between *intentionality* and *reflexivity* could lead to a more accurate explanation of some of their behaviors.

Now, can bees entertain states that exhibit what I called *internalized reflexive characters*, i.e., reflexive characters that are projected toward the system’s own mentality? There does not seem to be any evidence in support of such hypothesis. Moreover, since bees do not seem to be able to form individual representations, it is highly unlikely that they can individuate themselves from others. Consider, for instance, the fact that there usually are tens of thousands of bees living in the same hive: what would be the evolutionary advantage of the capacity to identify individuals? Certainly, we require more research to decide whether bees can entertain internal states that exhibit internalized reflexive characters, but the

observations reviewed so far suggest that we could explain their behaviors in terms of *intentionality* and, in some cases, in terms of a relation between *intentionality* and *reflexivity*. What seems to be clear is that the approach of traditional ethology is insufficient and, more importantly, that it may not be the simpler explanation of bee behavior.

4.2.2. Bee emotions?

There is no objective way to evaluate whether a system can entertain phenomenal states because the phenomenality of a conscious state cannot be identified from a third-person perspective¹¹⁵. Although theorists have been able to pinpoint neural structures and patterns of behavior that seem to be directly correlated to the phenomenal characters of human conscious states, the relation between consciousness and its basis, i.e., awareness, is still conceptually contingent.

So far, I have employed a one-dimensional model of consciousness based on *intentionality* to argue that some bee behaviors are best explained by describing their corresponding internal states in terms of *intentionality* and, in some cases, in terms of a relation between *intentionality* and *reflexivity*. However, I have not yet said anything about the possibility of attributing phenomenality to at least some bee internal states. In my view, since the components of *consciousness* are not independent from each other, we might already have reasons to consider that some bee internal states may exhibit phenomenal characters. For instance, given that bees can make perceptual distinctions and learn to modify their behavior according to context-dependent changes, it is plausible that their perceptual states are like something to them.

Arguably, a system capable of entertaining perceptual states is a system capable of entertaining states that exhibit phenomenal characters, insofar as perceptual features tend to be associated to phenomenal features¹¹⁶. The fact that bees can distinguish colors, patterns, shapes, etc., to modify their behavior suggests that they may be able to distinguish phenomenal features; after all, perceptions of colors and shapes are often said to have phenomenal features. If so, we might hypothesize that bees can entertain internal states that exhibit phenomenal characters. Hence, the question is whether they can identify and discriminate phenomenal features.

¹¹⁵ Notice that we can only make the phenomenality of our conscious states accessible to others by means of conceptualizing it according to a convention and expressing it linguistically, but this indirect method is not available when studying non-human systems.

¹¹⁶ Perhaps these are two different ways to describe the same features.

Theorists have recently suggested that bees exhibit emotional responses to a variety of stimuli. Broadly speaking, the idea is that “emotions can be identified as internal states that alter “cognitive” behavior, such as decision-making under conditions where the outcome of the decision is ambiguous” (Adolphs & Anderson, 2018, p. 210). Thus, if bees can modify their behavior according to the alleged phenomenality of their internal states, we might have a case to argue that they can entertain internal states that exhibit phenomenal characters. For instance, theorists have observed that “bees are more likely to classify ambiguous stimuli as predicting punishment” (Bateson et al., 2011, p. 1070), which could be a sign that bees can entertain states that have phenomenality, insofar as the behavior exhibited could be described in terms of the ability to identify and discriminate phenomenal features. However, such behaviors could also be explained as adaptive (Giurfa, 2013, pp. 288–289), in which case we would have no need to describe the system’s internal state in terms of *phenomenality*. Similarly, the emotion-like behavior exhibited by bumblebees when they receive an unexpected reward (Perry et al., 2016) could be explained as an increase of appetite (Núñez & Giurfa, 1996).

A popular approach to phenomenality is to evaluate whether a certain kind of system is capable of feeling pain. Broadly speaking, when we talk about pain, we talk about a certain kind of phenomenal character. Hence the common distinction between nociception, i.e., “the capacity to respond to potentially damaging stimuli” (Adamo, 2016, p. 75), and pain. However, this distinction suggests that the phenomenality of pain can be detached from its functional basis, which is a matter of debate. Since my proposal mainly concerns conceptual aspects of the study of consciousness, I do not need to take a stand concerning this debate. In my view, the study of consciousness belongs to an irreducible conceptual level, which does not mean that consciousness must be independent from its basis.

According to my interpretation of *consciousness*, the phenomenality of conscious states could be identified and conceptualized by relying on the interrelations between the components of *consciousness*. In the case of bees, the one-dimensional model of consciousness that I have been employing provides a means to investigate whether bees could be able to entertain internal states that exhibit phenomenal characters, insofar as their alleged phenomenality could be described in terms of how it relates to the intentionality of some of their internal states. Accordingly, while there seems to be no doubt that insects have nociception, it must also be acknowledged that they “show some differences in their responses to nociception compared to vertebrates. For example, insects tend to continue to

use damaged limbs [...], will eat their own innards [...] and will continue to feed while being consumed by another insect” (Adamo, 2016, p. 76). What these observations suggest is that, although pain must always be painful, different kinds of systems might experience it differently. Hence,

[...] the observation that insect behaviour differs from human behaviour when exposed to noxious stimuli does not necessarily mean that they do not have a pain-like experience. Being able to experience the emotional component of pain may not be an all-or-none phenomenon. Insects could have some aspects of an emotional experience but still lack the full subjective experience [...]. Moreover, this capacity may vary across species, depending on the whether or not a subjective experience of pain would provide a fitness advantage (Adamo, 2016, p. 76).

I do not know what a “full subjective experience” is, nor what a partial subjective experience would be. In any case, it seems that the issue could only be decided by establishing a frame of reference, like human consciousness, which would imply a description of bee consciousness in terms of *human consciousness*. Certainly, this is not what I have been proposing here. In my view, we should start by considering the last claim in the passage above, i.e., that the ability to experience pain must provide a “fitness advantage”, because there is no universal environment to which all systems must adapt.

Although we cannot know for sure whether bees are able to experience pain, it seems that they do not respond mechanically to noxious stimuli. Thus, a more accurate explanation of their behavior could be given in terms of a relation between *intentionality* and *phenomenality*, where their corresponding internal states could be described as states that target specific phenomenal features. For instance, despite the fact that insects behave very differently from us when their bodies are damaged, we could explain the complexity of their behavior by claiming that they might be able to identify phenomenal features and modify their behavior accordingly. What seems to be clear is that insects, in general, modify their behavior according to factors, like losing a limb, that only seem to be describable in terms of *phenomenality* because, otherwise, they would continue to behave as if nothing had happened.

Certainly, I have not provided evidence in favor of the hypothesis that bees are able to entertain states that exhibit phenomenal characters, but that was not my aim. My goal was to argue that the complexity of some bee behaviors might be best explained by describing their corresponding internal states in terms of relations between *intentionality* and *phenomenality*. Although more research is needed, what seems to be the case is that bees

behave in ways that do not seem to be reducible to mere mechanisms, which is why I have suggested that a more accurate explanation of bee behavior could be given by describing their internal states in terms of *consciousness*.

4.2.3. Can we attribute consciousness to bees?

In this section, I have analyzed observations of bee behavior in terms of my interpretation of *consciousness*. In particular, I have argued that the complexity of bee behavior could be explained in terms of a one-dimensional model of consciousness based on *intentionality*. My analysis suggests that some bee behaviors may be explained in terms of *consciousness*, which motivates my claim that the attribution of at least some aspects of consciousness to bees could provide a more accurate explanation of their activities. Of course, my analysis does not imply that bees are conscious systems, and the observations that I have reviewed do not evince that they are able to entertain conscious states. What my proposal implies is that an explanation of some bee behaviors in terms of *consciousness* could lead to a better understanding of why they behave as they do.

While it is reasonable to assume that non-human systems might entertain conscious states that are significantly different from human conscious states, it should also be acknowledged that every approach to non-human systems must be made from the human point of view; after all, the study of consciousness is a human endeavor. However, although certain patterns of human behavior are often interpreted as signs of consciousness, the relation between behavior and consciousness is conceptually contingent. Thus, the observations that I have analyzed suggest that, if bees are indeed able to entertain conscious states, those states can only be described by accounting for the way in which they interact with their environment and in terms of their sensory/cognitive skills, not in terms of *human consciousness*. Hence, my main claim is that, although we cannot be sure of whether bees entertain conscious states, an explanation of their behavior in terms of *consciousness* is simpler than what traditional ethology offers.

My analysis of bee behavior suggests that, although “we have no guarantee that animals that behave like humans do so because they have the same subjective experiences that humans do” (Barron & Klein, 2016, p. 4900), the fact that bees behave differently from humans does not mean that they are unable to entertain conscious states. Accordingly, the purpose of this section was to show how my interpretation of *consciousness* can be employed to study non-human systems whose environment is significantly different from ours, as well as their sensory/cognitive skills. Consequently, if we have good reasons to describe some

bee internal states in terms of *consciousness*, it might turn out that consciousness is more widespread than we normally think.

By means of a one-dimensional model of consciousness based on *intentionality*, I have argued that some bee behaviors can be explained by attributing intentional states to them and that those states might also exhibit reflexive and phenomenal characters. Accordingly, the idea is that the reflexivity and the phenomenality of their alleged conscious states must be described in terms of how they relate to the intentionality of their corresponding internal states. However, we must grant that, if bees entertain intentional states, those states cannot be compared to human intentional states, because their sensory/cognitive skills are significantly different from ours. Hence, the possibility to attribute consciousness to bees is a practice of conceptualization that is mediated by a proper interpretation of the way in which they interact with their environment, and not just a methodological resource to bypass the limits of traditional ethology.

Regarding reflexivity, we do not seem to have enough reasons to consider that bees could be able to entertain reflexive states, understood as conscious states that are mainly described in terms of *reflexivity*. Likewise, there do not seem to be enough reasons to think that bees can entertain phenomenal states, i.e., conscious states that are mainly described in terms of *phenomenality*. However, according to my analysis, the apparent intentionality of some of their internal states may have to be described in terms of relations between *intentionality* and the other two components of *consciousness*. Thus, it seems that attributing consciousness to bees could provide a more accurate understanding of their activities.

Of course, some might argue that a one-dimensional model of consciousness based on *intentionality* is enough to investigate non-human consciousness. Others could even suggest that the concept of *unconscious intentionality* would equally do the trick. Indeed, those would be valid pragmatic choices, legitimate theoretical strategies. However, the richness of dolphin and bee behavior that I have analyzed here at least suggests that consciousness might be a more widespread phenomenon than the latter strategies would suggest. The attribution of consciousness, in my view, is a practice of conceptualization and, as such, its purpose has to be made explicit. Otherwise, consciousness would seem to be no more than a human invention, in which case, consciousness would be *consciousness*.

Concluding Remarks

I have presented an interpretation of the concept of consciousness that does not rely on any particular metaphysical view concerning its ultimate nature. My goal was to argue that the study of consciousness belongs to an irreducible conceptual level and that, consequently, it must be explained in its own terms. Of course, this does not mean that consciousness is independent from, say, the mechanisms that enable it, its origin, or the role it plays. What my approach suggests is that, if consciousness is to be considered a legitimate subject-matter of scientific research, then we must explicate its concept. My approach to consciousness had two main purposes:

1. It is meant to be metaphysically compatible with most theories of consciousness and views about the ultimately correct analysis of its concept, insofar as it provides a conceptual background with which those theories and view can be assessed. By analyzing our uses of the term ‘consciousness’ and how theorists seem to characterize the concept of consciousness, I have argued that our characterization of the concept depends on the identification of three paradigmatic features exhibited by conscious states, namely, intentionality, reflexivity and phenomenality. Accordingly, I have introduced an interpretation of the concept of consciousness that is grounded on the thesis that our conceptualizations of those features, which I called *components* of the concept of consciousness, are interrelated, and that they form a three-dimensional conceptual space. Such conceptual space, I argued, constitutes the concept of consciousness. Since this interpretation of the concept is entirely based on considerations that concern our use of the

concept, it enables the assessment of different theories of consciousness from a single characterization of the concept, instead of a family or hierarchy of concepts of consciousness.

2. It motivates a different kind of theory of consciousness. Instead of providing reasons to endorse a specific view regarding the ultimate nature of consciousness, I have argued that consciousness can be explained by means of characterizing conscious states as descriptions of some of the internal states of a system in terms of the components of the concept of consciousness. What distinguishes my approach from most theories of consciousness is that it targets consciousness without determining its ontological status or providing a definite analysis of its concept. This is the reason why I have focused on how consciousness is attributed, not on what it is. In my view, the study of consciousness is motivated by our attributions of consciousness, which are always mediated by the implementation of what I called *practices of conceptualization*, not by direct observation.

What I have presented here is not a theory of consciousness, but a conceptual framework meant to explicate the concept of consciousness. A theory of consciousness should explain the phenomena to which we refer when we apply the concept of consciousness, and my proposal only concerns that concept. However, in the third chapter, I introduced a characterization of conscious states that could be interpreted as the basic structure of a theory of consciousness, inasmuch as, along with the interpretation of the concept of consciousness that I presented in the second chapter, it allows us to theorize about possible cases of non-human consciousness, as I showed in the fourth chapter.

Certainly, my approach does not answer any of the problems that constitute the contemporary debate on consciousness, though it suggests some ways to address them. For instance, concerning the so-called *hard problem of consciousness*, my view indicates that the reason why an answer to the so-called *easy problems* would not lead to an answer to the hard one is that the latter belongs to a different conceptual level. According to my analysis in the first chapter, the easy problems concern the functional basis of consciousness, namely, awareness, and I have argued that an explanation of consciousness cannot be reduced to an explanation of awareness because an explanation of the mechanisms that allegedly underlie consciousness would not explain the nature of experience. However, my approach is silent with regards to the problems that concern the metaphysics of consciousness, inasmuch as it is equally compatible with the claims that consciousness exists, that it is reducible to something else, and that it does not exist. My proposal is grounded on the fact that we

attribute consciousness, not on what it is because this, as I have acknowledged since the beginning of this work, I do not know.

One advantage of my approach is that it enables a sort of taxonomy of consciousness that does not rely on how we characterize human consciousness. Indeed, the study of consciousness is a human endeavor, but that does not mean that every possible kind of conscious system must entertain a human-like kind of consciousness. In my view, the attribution of consciousness is based on our ability to identify and conceptualize certain paradigmatic characters exhibited by conscious states, not on the specific ways in which those characters are exhibited by human conscious states. Thus, my approach allows us to contrast kinds of conscious systems in terms of a single concept, not in reference to the kind of consciousness that we attribute to ourselves.

Another advantage of my view is that it turns consciousness into a legitimate subject-matter of scientific research. Our attributions of consciousness suggest that consciousness is part of our world, even if it were a mere illusion. The fact that the attribution of consciousness is a practice of conceptualization that concerns the way in which we interpret the world means that it is part of our worldview, regardless of whether there are conscious systems or not in the metaphysical sense. Thus, the study of consciousness is legitimized by the fact that it plays a role in the construction of our worldview.

In summary, what I have presented here is an approach to the study of consciousness that is based on how we understand its concept. What motivated my view is the observation that, although there are outstanding theories of consciousness, there is no clear way to evaluate them or decide which one of them best explains it. Consciousness might be, as many theorists claim, a phenomenon with which we are very well acquainted, but that does not mean that it can be explained by the standard methods of contemporary science. In my opinion, explaining consciousness does not depend on providing irrefutable arguments in favor of such-and-such theory or evidence of its existence, for the soundness of an argument depends on the practice of conceptualization adopted, as well as the sufficiency of evidence. Consciousness is not what needs an explanation, we need to explain it.

Bibliography

- Adamo, S. (2016). Do insects feel pain? A question at the intersection of animal behaviour, philosophy and robotics. *Animal Behaviour*, *118*, 75–79.
<https://doi.org/10.1016/j.anbehav.2016.05.005>
- Adolphs, R., & Anderson, D. (2018). Emotions in Insects and Other Invertebrates. In *The Neuroscience of Emotion* (pp. 197–214). Princeton University Press.
<https://doi.org/10.2307/j.ctvc77b1j.12>
- Alem, S., Perry, C., Zhu, X., Loukola, O., Ingraham, T., Søvik, E., & Chittka, L. (2016). Associative Mechanisms Allow for Social Learning and Cultural Transmission of String Pulling in an Insect. *PLOS Biology*, *14*(10).
<https://doi.org/10.1371/journal.pbio.1002564>
- Allen, C., & Bekoff, M. (1999). *Species of mind: The philosophy and biology of cognitive ethology*. MIT.
- Avarguès-Weber, A., Dyer, A., Combe, M., & Giurfa, M. (2012). Simultaneous mastering of two abstract concepts by the miniature brain of bees. *Proceedings of the National Academy of Sciences*, *109*(19), 7481–7486.
<https://doi.org/10.1073/pnas.1202576109>
- Avarguès-Weber, A., Dyer, A., & Giurfa, M. (2011). Conceptualization of above and below relationships by an insect. *Proceedings of the Royal Society B: Biological Sciences*, *278*(1707), 898–905. <https://doi.org/10.1098/rspb.2010.1891>
- Baars, B. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
- Baars, B. (2005). Subjective experience is probably not limited to humans: The evidence from neurobiology and behavior. *Consciousness and Cognition*, *14*(1), 7–21.
<https://doi.org/10.1016/j.concog.2004.11.002>
- Barron, A., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(18), 4900–4908.

- Bateson, M., Desire, S., Gartside, S., & Wright, G. (2011). Agitated Honeybees Exhibit Pessimistic Cognitive Biases. *Current Biology*, 21(12), 1070–1073. <https://doi.org/10.1016/j.cub.2011.05.017>
- Bayne, T., Cleeremans, A., & Wilken, P. (Eds.). (2009). *The Oxford companion to consciousness*. Oxford University Press.
- Bayne, T., Hohwy, J., & Owen, A. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences*, 20(6), 405–413. <https://doi.org/10.1016/j.tics.2016.03.009>
- Bermúdez, J. L. (1998). *The paradox of self-consciousness*. MIT Press.
- Bermúdez, J. L. (2011). *Bodily Awareness and Self-Consciousness*. Oxford University Press. <http://oxfordhandbooks.com/view/10.1093/oxfordhb/9780199548019.001.0001/oxfordhb-9780199548019-e-7>
- Blackburn, S. (1996). *The Oxford dictionary of philosophy*. Oxford University Press.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227–247. <https://doi.org/10.1017/S0140525X00038188>
- Bourget, D., & Mendelovici, A. (2019). Phenomenal Intentionality. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University. <https://plato.stanford.edu/archives/fall2019/entries/phenomenal-intentionality/>
- Brentano, F. (2005). *Psychology from an empirical standpoint* (O. Kraus & L. McAlister, Eds.). Routledge. <http://www.dawsonera.com/depp/reader/protected/external/AbstractView/S9780203202883>
- Brüntrup, G., & Jaskolla, L. (Eds.). (2017). *Panpsychism: Contemporary perspectives*. Oxford University Press.
- Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press.
- Carnap, R. (2003). *The logical structure of the world: And, pseudoproblems in philosophy*. Open Court.
- Carruthers, P. (2003). *Phenomenal consciousness: A naturalistic theory*. Cambridge University Press.
- Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. Oxford University Press.
- Chalmers, D., & Bayne, T. (2003). What is the unity of consciousness? In A. Cleeremans & C. Frith (Eds.), *The Unity of Consciousness: Binding, Integration, and Dissociation* (pp. 23–58). Oxford University Press.
- Chittka, L., Williams, N., Rasmussen, H., & Thomson, J. (1999). Navigation without Vision: Bumblebee Orientation in Complete Darkness. *Proceedings: Biological Sciences*, 266(1414), 45–50.
- Churchland, P. (1995). *The engine of reason, the seat of the soul: A philosophical journey into the brain*. MIT Press.

- Coleman, S. (2012). Mental Chemistry: Combination for Panpsychists. *Dialectica*, 66(1), 137–166.
- Coleman, S. (2015a). Neuro-Cosmology. In P. Coates & S. Coleman (Eds.), *Phenomenal Qualities* (pp. 66–102). Oxford University Press.
<https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780198712718.001.0001/acprof-9780198712718-chapter-3>
- Coleman, S. (2015b). Quotational higher-order thought theory. *Philosophical Studies*, 172(10), 2705–2733. <https://doi.org/10.1007/s11098-015-0441-1>
- Crane, T. (2001). *Elements of mind: An introduction to the philosophy of mind*. Oxford University Press.
- Crane, T. (2013). *The objects of thought*. Oxford University Press.
- Daston, L., & Mitman, G. (2007). *Thinking with animals: New perspectives on anthropomorphism*. Columbia University Press.
- de Vignemont, F. (2018). Bodily Awareness. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
<https://plato.stanford.edu/archives/spr2018/entries/bodily-awareness/>
- Dennett, D. (1971). Intentional Systems. *The Journal of Philosophy*, 68(4), 87–106.
<https://doi.org/10.2307/2025382>
- Dennett, D. (1991). *Consciousness explained*. Little, Brown and Co.
- Dennett, D. (1995). Animal Consciousness: What Matters and Why. *Social Research*, 62(3), 691–710.
- Descartes, R. (1984). Meditations on First Philosophy. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The Philosophical Writings of Descartes* (Vol. 2, pp. 1–397). Cambridge University Press.
- Dewey, J. (1906). The Terms “Conscious” and “Consciousness.” *The Journal of Philosophy, Psychology and Scientific Methods*, 3(2), 39–41.
<https://doi.org/10.2307/2011714>
- Fasching, W. (2009). The mineness of experience. *Continental Philosophy Review*, 42(2), 131–148. <https://doi.org/10.1007/s11007-009-9107-z>
- Flanagan, O. (1995). *Consciousness reconsidered*. MIT Press.
- Gennaro, R. (2012). *The consciousness paradox: Consciousness, concepts, and higher-order thoughts*. MIT Press.
- Gertler, B. (2019). Self-Knowledge. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
<https://plato.stanford.edu/archives/win2019/entries/self-knowledge/>
- Giger, A., & Srinivasan, M. (1995). Pattern recognition in honeybees: Eidetic imagery and orientation discrimination. *Journal of Comparative Physiology A*, 176(6), 791–795.
<https://doi.org/10.1007/BF00192626>
- Giurfa, M. (2013). Cognition with few neurons: Higher-order learning in insects. *Trends in Neurosciences*, 36(5), 285–294. <https://doi.org/10.1016/j.tins.2012.12.011>
- Giurfa, M., & Lehrer, M. (2005). Honeybee vision and floral displays: From detection to close-up recognition. In L. Chittka & J. Thomson (Eds.), *Cognitive ecology of*

- pollination: Animal behavior and floral evolution* (pp. 61–82). Cambridge University Press.
- Giurfa, M., Núñez, J., & Backhaus, W. (1994). Odour and colour information in the foraging choice behaviour of the honeybee. *Journal of Comparative Physiology A*, *175*(6), 773–779. <https://doi.org/10.1007/BF00191849>
- Giurfa, M., Vorobyev, M., Brandt, R., Posner, B., & Menzel, R. (1997). Discrimination of coloured stimuli by honeybees: Alternative use of achromatic and chromatic signals. *Journal of Comparative Physiology A*, *180*(3), 235–243. <https://doi.org/10.1007/s003590050044>
- Giurfa, M., Vorobyev, M., Kevan, P., & Menzel, R. (1996). Detection of coloured stimuli by honeybees: Minimum visual angles and receptor specific contrasts. *Journal of Comparative Physiology A*, *178*(5), 699–709. <https://doi.org/10.1007/BF00227381>
- Giurfa, M., Zhang, S., Jenett, A., Menzel, R., & Srinivasan, M. (2001). The concepts of ‘sameness’ and ‘difference’ in an insect. *Nature*, *410*(6831), 930–933. <https://doi.org/10.1038/35073582>
- Goff, P. (2017). *Consciousness and fundamental reality*. Oxford University Press.
- Guillot, M., & Garcia-Carpintero, M. (forthcoming). *The Sense of Mineness*. Oxford University Press.
- Hanke, W., & Erdsack, N. (2015). Ecology and Evolution of Dolphin Sensory Systems. In D. Herzing & C. Johnson (Eds.), *Dolphin communication and cognition: Past, present, and future* (pp. 49–74). The MIT Press.
- Harley, H. (2013). Consciousness in dolphins? A review of recent evidence. *Journal of Comparative Physiology*, *199*(6), 565–582. <https://doi.org/10.1007/s00359-013-0816-8>
- Harley, H., Roitblat, H., & Nachtigall, P. (1996). Object representation in the bottlenose dolphin (*Tursiops truncatus*): Integration of visual and echoic information. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*(2), 164–174. <https://doi.org/10.1037/0097-7403.22.2.164>
- Herman, L. (2002a). Exploring the Cognitive World of the Bottlenosed Dolphin. In M. Bekoff, C. Allen, & G. Burghardt (Eds.), *The cognitive animal: Empirical and theoretical perspectives on animal cognition* (pp. 275–283). MIT Press.
- Herman, L. (2002b). Vocal, social, and self-imitation by bottlenosed dolphins. In *Imitation in animals and artifacts* (pp. 63–108). Boston Review.
- Herman, L. (2006). Intelligence and rational behaviour in the bottlenosed dolphin. In S. Hurley & M. Nudds (Eds.), *Rational Animals?* (pp. 439–468). Oxford University Press. <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780198528272.001.0001/acprof-9780198528272-chapter-20>
- Herman, L. (2012). Body and self in dolphins. *Consciousness and Cognition*, *21*(1), 526–545. <https://doi.org/10.1016/j.concog.2011.10.005>
- Herman, L., & Forestell, P. (1985). Reporting presence or absence of named objects by a language-trained dolphin. *Neuroscience & Biobehavioral Reviews*, *9*(4), 667–681. [https://doi.org/10.1016/0149-7634\(85\)90013-2](https://doi.org/10.1016/0149-7634(85)90013-2)

- Herman, L., & Gordon, J. (1974). Auditory delayed matching in the bottlenose dolphin. *Journal of the Experimental Analysis of Behavior*, 21(1), 19–26. <https://doi.org/10.1901/jeab.1974.21-19>
- Herman, L., Hovancik, J., Gory, J., & Bradshaw, G. (1989). Generalization of visual matching by a bottlenosed dolphin (*Tursiops truncatus*): Evidence for invariance of cognitive performance with visual and auditory materials. *Journal of Experimental Psychology: Animal Behavior Processes*, 15(2), 124–136. <https://doi.org/10.1037/0097-7403.15.2.124>
- Herman, L., Kuczaj, S., & Holder, M. (1993). Responses to anomalous gestural sequences by a language-trained dolphin: Evidence for processing of semantic relations and syntactic information. *Journal of Experimental Psychology: General*, 122(2), 184–194. <https://doi.org/10.1037/0096-3445.122.2.184>
- Herman, L., Matus, D., Herman, E., Ivancic, M., & Pack, A. (2001). The bottlenosed dolphin's (*Tursiops truncatus*) understanding of gestures as symbolic representations of its body parts. *Animal Learning & Behavior*, 29(3), 250–264. <https://doi.org/10.3758/BF03192891>
- Herman, L., Pack, A., & Hoffmann-Kuhnt, M. (1998). Seeing through sound: Dolphins (*Tursiops truncatus*) perceive the spatial structure of objects through echolocation. *Journal of Comparative Psychology*, 112(3), 292–305. <https://doi.org/10.1037/0735-7036.112.3.292>
- Herzing, D., & Johnson, C. (Eds.). (2015). *Dolphin communication and cognition: Past, present, and future*. The MIT Press.
- Horridge, A. (2005). What the honeybee sees: A review of the recognition system of *Apis mellifera*. *Physiological Entomology*, 30(1), 2–13. <https://doi.org/10.1111/j.0307-6962.2005.00425.x>
- Houdé, O. (Ed.). (2004). *Dictionary of cognitive science: Neuroscience, psychology, artificial intelligence, linguistics, and philosophy*. Psychology Press.
- Jackson, F. (1986). What Mary Didn't Know. *The Journal of Philosophy*, 83(5), 291–295. <https://doi.org/10.2307/2026143>
- Jacob, P. (2019). Intentionality. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University. <https://plato.stanford.edu/archives/win2019/entries/intentionality/>
- Janik, V. (2013). Cognitive skills in bottlenose dolphin communication. *Trends in Cognitive Sciences*, 17(4), 157–159. <https://doi.org/10.1016/j.tics.2013.02.005>
- Johnson, C., Sullivan, J., Buck, C., Trexel, J., & Scarpuzzi, M. (2015). Visible and invisible displacement with dynamic visual occlusion in bottlenose dolphins (*Tursiops* spp). *Animal Cognition*, 18(1), 179–193. <https://doi.org/10.1007/s10071-014-0788-2>
- Kriegel, U. (2009). *Subjective consciousness: A self-representational theory*. Oxford University Press.
- Kuczaj, S., & Winship, K. (2015). How Do Dolphin Calves Make Sense of Their World? In D. Herzing & C. Johnson (Eds.), *Dolphin communication and cognition: Past, present, and future* (pp. 201–226). The MIT Press.

- Lammers, M., & Oswald, J. (2015). Analyzing the Acoustic Communication of Dolphins. In D. Herzing & C. Johnson (Eds.), *Dolphin communication and cognition: Past, present, and future* (pp. 107–137). The MIT Press.
- Levine, J. (1983). Materialism and Qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64(4), 354–361. <https://doi.org/10.1111/j.1468-0114.1983.tb00207.x>
- Locke, J. (1975). *An essay concerning human understanding* (P. Nidditch, Ed.). Oxford University Press.
- Loveland, K. (1995). Self-Recognition in the Bottlenose Dolphin: Ecological Considerations. *Consciousness and Cognition*, 4(2), 254–257. <https://doi.org/10.1006/ccog.1995.1035>
- Lycan, W. (1996). *Consciousness and experience*. MIT Press.
- Maffie, J. (2008). Consciousness and Reality in Nahua Thought in the Era of the Conquest. In H. Wautischer (Ed.), *Ontology of consciousness: Percipient action* (pp. 97–126). MIT Press.
- Marino, L., Connor, R., Fordyce, R. E., Herman, L., Hof, P., Lefebvre, L., Lusseau, D., McCowan, B., Nimchinsky, E., Pack, A., Rendell, L., Reidenberg, J., Reiss, D., Uhen, M., Van der Gucht, E., & Whitehead, H. (2007). Cetaceans Have Complex Brains for Complex Cognition. *PLoS Biology*, 5(5). <https://doi.org/10.1371/journal.pbio.0050139>
- Márquez Velasco, L. (2013). La razón racionalista: Exploración de un espacio conceptual. *Episteme NS*, 33(1), 99–115.
- McGinn, C. (1991a). Consciousness and Content. In R. Bogdan (Ed.), *Mind and Common Sense* (pp. 71–92). Cambridge University Press. https://www.cambridge.org/core/product/identifier/CBO9780511551659A010/type/book_part
- McGinn, C. (1991b). *The problem of consciousness: Essays towards a resolution*. Blackwell.
- Mendelovici, A. (2018). *The phenomenal basis of intentionality*. Oxford University Press.
- Menzel, R., & Giurfa, M. (2001). Cognitive architecture of a mini-brain: The honeybee. *Trends in Cognitive Sciences*, 5(2), 62–71. [https://doi.org/10.1016/S1364-6613\(00\)01601-6](https://doi.org/10.1016/S1364-6613(00)01601-6)
- Miyamoto, Y., Nisbett, R., & Masuda, T. (2006). Culture and the Physical Environment: Holistic versus Analytic Perceptual Affordances. *Psychological Science*, 17(2), 113–119.
- Morgan, C. L. (1903). *An introduction to comparative psychology*. Walter Scott Publishing Co. <http://content.apa.org/books/13701-000>
- Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
- Núñez, J., & Giurfa, M. (1996). Motivation and regulation of honey bee foraging. *Bee World*, 77(4), 182–196. <https://doi.org/10.1080/0005772X.1996.11099313>

- Pack, A. (2015). Experimental Studies of Dolphin Cognitive Abilities. In D. Herzing & C. Johnson (Eds.), *Dolphin communication and cognition: Past, present, and future* (pp. 175–200). The MIT Press.
- Pack, A., & Herman, L. (1995). Sensory integration in the bottlenosed dolphin: Immediate recognition of complex shapes across the senses of echolocation and vision. *The Journal of the Acoustical Society of America*, *98*(2), 722–733. <https://doi.org/10.1121/1.413566>
- Peacocke, C. (2014). *The mirror of the world: Subjects, consciousness, and self-consciousness*. Oxford University Press.
- Perry, C., Baciadonna, L., & Chittka, L. (2016). Unexpected rewards induce dopamine-dependent positive emotion-like state changes in bumblebees. *Science*, *353*(6307), 1529–1531. <https://doi.org/10.1126/science.aaf4454>
- Perry, C., & Barron, A. (2013). Honey bees selectively avoid difficult choices. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(47), 19155–19159.
- Pinker, S. (2009). *How the Mind Works*. W.W. Norton & Company.
- Pitt, D. (2004). The Phenomenology of Cognition Or What Is It Like to Think That P? *Philosophy and Phenomenological Research*, *69*(1), 1–36. <https://doi.org/10.1111/j.1933-1592.2004.tb00382.x>
- Putnam, H. (1967). Psychological Predicates. In W. Capitan & D. Merrill (Eds.), *Art, Mind, and Religion* (pp. 37–48). University of Pittsburgh Press.
- Pylyshyn, Z. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, *80*(1–2), 127–158. [https://doi.org/10.1016/S0010-0277\(00\)00156-6](https://doi.org/10.1016/S0010-0277(00)00156-6)
- Reiss, D., & Marino, L. (2001). Mirror Self-Recognition in the Bottlenose Dolphin: A Case of Cognitive Convergence. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(10), 5937–5942.
- Rendell, L., & Whitehead, H. (2001). Culture in whales and dolphins. *Behavioral and Brain Sciences*, *24*(2), 309–324. <https://doi.org/10.1017/S0140525X0100396X>
- Ridgway, S., Carder, D., Finneran, J., Keogh, M., Kamolnick, T., Todd, M., & Goldblatt, A. (2006). Dolphin continuous auditory vigilance for five days. *Journal of Experimental Biology*, *209*, 3621–3628. <https://doi.org/10.1242/jeb.02405>
- Roelofs, L. (2014). Phenomenal Blending and the Palette Problem. *Thought: A Journal of Philosophy*, *3*(1), 59–70. <https://doi.org/10.1002/tht3.113>
- Roelofs, L. (2016). The unity of consciousness, within subjects and between subjects. *Philosophical Studies*, *173*(12), 3199–3221. <https://doi.org/10.1007/s11098-016-0658-7>
- Roitblat, H. (2002). The Cognitive Dolphin. In M. Bekoff, C. Allen, & G. Burghardt (Eds.), *The cognitive animal: Empirical and theoretical perspectives on animal cognition* (pp. 183–188). MIT Press.
- Rosenthal, D. (1993). Thinking that one thinks. In M. Davies & G. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 197–223). Blackwell.
- Searle, J. (1992). *The rediscovery of the mind*. MIT Press.

- Searle, J. (1994). Animal Minds. *Midwest Studies In Philosophy*, 19(1), 206–219.
<https://doi.org/10.1111/j.1475-4975.1994.tb00286.x>
- Searle, J. (2017). Biological Naturalism. In S. Schneider & M. Velmans (Eds.), *The Blackwell companion to consciousness* (pp. 327–336). John Wiley & Sons Inc.
- Sheldrake, M. (2020). *Entangled life: How fungi make our worlds, change our minds & shape our futures*. Random House.
- Siderits, M., Thompson, E., & Zahavi, D. (2010). *Self, No Self?* Oxford University Press.
<http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199593804.001.0001/acprof-9780199593804>
- Siewert, C. (2016). Consciousness and Intentionality. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
<https://plato.stanford.edu/archives/spr2017/entries/consciousness-intentionality/>
- Skorupski, P., MaBouDi, H., Galpayage Dona, H., & Chittka, L. (2018). Counting insects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1740).
<https://doi.org/10.1098/rstb.2016.0513>
- Slors, M., & Jongepier, F. (2014). Mineness without Minimal Selves. *Journal of Consciousness Studies*, 21(7–8), 193–219.
- Smith, J. (2017). Self-Consciousness. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
<https://plato.stanford.edu/archives/fall2017/entries/self-consciousness/>
- Solvi, C., Gutierrez Al-Khudhairy, S., & Chittka, L. (2020). Bumble bees display cross-modal object recognition between visual and tactile senses. *Science*, 367(6480), 910–912. <https://doi.org/10.1126/science.aay8064>
- Srinivasan, M. (2010). Honey Bees as a Model for Vision, Perception, and Cognition. *Annual Review of Entomology*, 55(1), 267–284.
<https://doi.org/10.1146/annurev.ento.010908.164537>
- Stamp Dawkins, M. (1998). *Through our eyes only? The search for animal consciousness*. Oxford University Press.
- Strawson, G. (2004). Real intentionality. *Phenomenology and the Cognitive Sciences*, 3(3), 287–313. <https://doi.org/10.1023/B:PHEN.0000049306.63185.0f>
- Strawson, G. (2017). Mind and Being: The Primacy of Panpsychism. In G. Brüntrup & L. Jaskolla (Eds.), *Panpsychism: Contemporary perspectives* (pp. 75–112). Oxford University Press.
- Strawson, P. (1992). *Analysis and metaphysics: An introduction to philosophy*. Oxford University Press.
- Sytsma, J., & Machery, E. (2009). How to study folk intuitions about phenomenal consciousness. *Philosophical Psychology*, 22(1), 21–35.
<https://doi.org/10.1080/09515080802703653>
- Tinker, G. (2004). The Stones Shall Cry out: Consciousness, Rocks, and Indians. *Wicazo Sa Review*, 19(2), 105–125.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. MIT Press.

- Van Gulick, R. (2018). Consciousness. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
<https://plato.stanford.edu/archives/spr2018/entries/consciousness>
- Velmans, M. (2009). *Understanding consciousness*. Routledge.
<http://site.ebrary.com/id/10289056>
- Wilson, R., & Keil, F. (1999). *The MIT encyclopedia of the cognitive sciences*. MIT Press.
- Wittlinger, M., Wehner, R., & Wolf, H. (2006). The Ant Odometer: Stepping on Stilts and Stumps. *Science*, 312(5782), 1965–1967. <https://doi.org/10.1126/science.1126912>
- Woodgate, J., Makinson, J., Lim, K., Reynolds, A., & Chittka, L. (2016). Life-Long Radar Tracking of Bumblebees. *PLOS ONE*, 11(8).
<https://doi.org/10.1371/journal.pone.0160333>
- Würsig, B., & Pearson, H. (2015). Dolphin Societies: Structure and Function. In D. Herzing & C. Johnson (Eds.), *Dolphin communication and cognition: Past, present, and future* (pp. 77–105). The MIT Press.
- Wystrach, A., & Graham, P. (2012). What can we learn from studies of insect navigation? *Animal Behaviour*, 84(1), 13–20. <https://doi.org/10.1016/j.anbehav.2012.04.017>
- Yaman, S., Kilian, A., von Fersen, L., & Güntürkün, O. (2012). Evidence for a Numerosity Category that is Based on Abstract Qualities of “Few” vs. “Many” in the Bottlenose Dolphin (*Tursiops truncatus*). *Frontiers in Psychology*, 3.
<https://doi.org/10.3389/fpsyg.2012.00473>
- Zahavi, D. (2000). Self and consciousness. In D. Zahavi (Ed.), *Exploring the Self. Philosophical and psychopathological perspectives on self-experience* (Vol. 23, pp. 55–74). John Benjamins Publishing Company.
<https://doi.org/10.1075/aicr.23.06zah>
- Zahavi, D. (2005). *Subjectivity and selfhood: Investigating the first-person perspective*. MIT Press.
- Zhang, S., & Srinivasan, M. (1994). Prior experience enhances pattern discrimination in insect vision. *Nature*, 368(6469), 330–332. <https://doi.org/10.1038/368330a0>