


EXPERIENCED WHOLENESS

INTEGRATING INSIGHTS FROM
GESTALT THEORY,
COGNITIVE NEUROSCIENCE,
AND PREDICTIVE PROCESSING

WANJA WIESE



Experienced Wholeness

Integrating Insights from Gestalt Theory, Cognitive Neuroscience, and Predictive Processing

Wanja Wiese

**The MIT Press
Cambridge, Massachusetts
London, England**

© 2017 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form or by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Stone Serif by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Wiese, Wanja, author.

Title: Experienced wholeness: integrating insights from Gestalt theory, cognitive neuroscience, and predictive processing / Wanja Wiese.

Description: Cambridge, MA: MIT Press, 2017. | Includes bibliographical references and index.

Identifiers: LCCN 2017015380 | ISBN 9780262036993 (hardcover: alk. paper)

Subjects: LCSH: Consciousness. | Phenomenology. | Gestalt psychology. | Cognitive neuroscience. | Inference.

Classification: LCC B808.9. W54 2017 | DDC 128/.2—dc23

LC record available at <https://lcn.loc.gov/2017015380>

To my parents, Ursula Wiese and Theodor Wiese

Contents

List of Figures and Boxes

Preface and Acknowledgments

I Setting the Scene

1 Introduction

- 1.1 A Brief Sketch of the Account Developed and Defended in This Book
- 1.2 Reflection on Methods
 - 1.2.1 Why an Interdisciplinary Account?
 - 1.2.2 How Is an Interdisciplinary Account of Phenomenal Unity Possible?
- 1.3 Phenomenal Unity and Access Unity
 - 1.3.1 What Aspects of Conscious Experience Are Accessible in the First Place?
 - 1.3.2 Access Unity
- 1.4 Why Anti-Representationalists Should Read This Book

2 Framing the Problem of Phenomenal Unity (PPU)

- 2.1 Basic Conceptual Considerations
 - 2.1.1 What Is a Phenomenal Field?
 - 2.1.2 What Is a Phenomenal Part?
 - 2.1.3 What Is a Subject of Experience?
 - 2.1.4 The First and the Second Problems of Phenomenal Unity
- 2.2 Diachronic and Synchronic Unity
- 2.3 Core Constraints on Accounts of Phenomenal Unity
 - 2.3.1 Constraint 1: Phenomenality
 - 2.3.2 Constraint 2: Globality
 - 2.3.3 Constraint 3: Necessity
- 2.4 Interim Conclusion

II The First Problem of Phenomenal Unity

3 Why the Single State Conception Does Not Solve the Problem of Phenomenal Unity

- 3.1 The Single State Conception (SSC)
- 3.2 SSC Creates Avoidable Problems
- 3.3 SSC Does Not Solve PPU
 - 3.3.1 Tim Bayne's Account of Phenomenal Unity Does Not Solve 1PPU
 - 3.3.2 The Metaphysical Digression Suggested by Bayne's Account Is a Dead End
- 3.4 Interim Conclusion

4 What Is It Like to Experience a Third Man? The Phenomenological Bradley and How to Solve It

- 4.1 The Phenomenological Bradley: A Phenomenological Dual to Bradley's Regress
- 4.2 Phenomenal Fundament
 - 4.2.1 Short Description
 - 4.2.2 Historical Precursors
- 4.3 Phenomenal Glue
 - 4.3.1 Short Description
 - 4.3.2 A Historical Precursor and One Current Account
- 4.4 Phenomenal Replacement
 - 4.4.1 Short Description
 - 4.4.2 Historical Precursors and One Current Account
- 4.5 Phenomenal Disposition
 - 4.5.1 Short Description
 - 4.5.2 One Historical Precursor and Two Current Accounts
- 4.6 The M \acute{e} lange Model: Integrating Different Solutions
- 4.7 Interim Conclusion

5 Think Globally, Experience Locally?

- 5.1 Understanding Globality
- 5.2 Holism and Consciousness

- 5.2.1 Michael Esfeld's Generic Holism
- 5.2.2 Holistic Theories of Consciousness
- 5.3 Strong Phenomenal Holism?
 - 5.3.1 Aristotle: Dynamic Unity
 - 5.3.2 Max Wertheimer: Principles of Perceptual Organization
 - 5.3.3 Aron Gurwitsch: Whole-Parts
 - 5.3.4 Elijah Chudnoff's Phenomenal Holism
 - 5.3.5 John Searle on Building Block and Unified Field Theories
 - 5.3.6 Tim Bayne on Atomistic and Holistic Theories of Consciousness
 - 5.3.7 Barry Dainton's Phenomenal Holism
- 5.4 The Plausible Alternative: Relax and Go Hierarchical
- 5.5 Multiplicity
 - 5.5.1 The Intuition
 - 5.5.2 Historical Precursors and Relations to Existing Accounts
 - 5.5.3 Summary
- 5.6 Graduality
 - 5.6.1 The Intuition
 - 5.6.2 Historical Precursors and Relations to Existing Accounts
 - 5.6.3 Summary
- 5.7 Invariance
 - 5.7.1 The Intuition
 - 5.7.2 Historical Precursors and Relations to Existing Accounts
 - 5.7.3 Summary
- 5.8 The Regularity Account of Phenomenal Unity
 - 5.8.1 An Important Precursor: Thomas Metzinger's Convolved Holism
 - 5.8.2 The Core Idea: Generalizing the Law of Common Fate
- 5.9 Interim Conclusion

III The Second Problem of Phenomenal Unity

6 Representational Holism

- 6.1 What Is the Representationalist Level of Description?
- 6.2 Global Self-Representations
 - 6.2.1 Robert van Gulick: Higher-Order Global States
 - 6.2.2 Uriah Kriegel: Cross-Order Integration
 - 6.2.3 Carolyn Jennings: Perceptual Organization
 - 6.2.4 Thomas Metzinger: What Is a Phenomenal Point of View?
- 6.3 What Are the Potential Benefits of a Hierarchy of Wholes?
 - 6.3.1 Types of Efficiency Associated with a Hierarchy of Wholes
 - 6.3.2 Segregation and Integration
- 6.4 Operationalizing Representational Holism: Causal Density, Integrated Information, and Related Measures
- 6.5 Interim Conclusion

7 Predictive Processing

- 7.1 Basics of Predictive Processing
- 7.2 PP and Karl Friston's Free-Energy Principle
- 7.3 PP and Bayesian Inference
- 7.4 PP and Active Inference
- 7.5 PP and Its Relation to Action and Perception
- 7.6 PP and Representationalism
 - 7.6.1 Are Hierarchical Generative Models in PP Genuine Representations?
 - 7.6.2 Ramsey's Job Description Challenge—Using Vehicles of Contents as Representations
 - 7.6.3 The Content Challenge
 - 7.6.4 What Is the Structural Concept of Representational Content?
 - 7.6.5 PP Entails a Structural Resemblance between Processes in the World and Processes in the Brain
 - 7.6.6 The Hierarchical Generative Model Plays a Representational Role
 - 7.6.7 How Are Particular Contents Determined?
- 7.7 PP and the Veil of Perception
- 7.8 Interim Conclusion

8 Consciousness and Predictive Processing

- 8.1 PP and Attention
 - 8.1.1 Predictive Weight Watching
 - 8.1.2 Endogenous and Exogenous Attention
 - 8.1.3 Is Volitional Attention Special?
 - 8.1.4 Volitional Attention and the Contents of Consciousness
 - 8.1.5 Precision Estimates and Affective Value
- 8.2 Jakob Hohwy on PP and Conscious Perception
- 8.3 PP and Theories of Consciousness
 - 8.3.1 Global Workspace Theories

8.3.2 Attention Schema Theory

8.3.3 Precision, the Global Workspace, and the Phenomenal Point of View

8.3.4 Integrated Information Theory

8.4 Interim Conclusion

9 The Regularity Account of Phenomenal Unity

9.1 Introducing RPU

9.1.1 Example 1: Auditory Regularities

9.1.2 Example 2: Visual Regularities

9.1.3 From Gestalt Laws to Predictive Regularities

9.2 What Is a Regularity?

9.2.1 An Informal Characterization

9.2.2 A Formal Characterization

9.2.3 A Possible Neural Implementation

9.2.4 A Hierarchy of Regularities

9.2.5 Regularities and INVARIANCE

9.2.6 The Role of Counterfactuals

9.3 A Hierarchy of Wholes

9.3.1 Feature Binding

9.3.2 Sensorimotor Integration

9.3.3 Multisensory Integration

9.3.4 Intramodal Perceptual Grouping

9.3.5 Intermodal Perceptual Grouping

9.3.6 Synesthesia

9.3.7 Crossmodal Correspondences

9.3.8 Emotions

9.3.9 Volitional Attention

9.4 A Formal Analysis of the Mélange Model

9.5 Which Constraints Are Satisfied by RPU?

9.5.1 Phenomenality

9.5.2 Globality

9.5.3 Necessity

9.6 Implications of RPU

9.6.1 States of Decreased Wholeness

9.6.2 States of Increased Wholeness

9.6.3 Attention Does Not Correlate with Wholeness

9.7 Interim Conclusion

10 Putting It All Together

10.1 A Selective Summary

10.2 Suggestions for Future Research

Glossary

References

Index

List of Figures

- 2.1 Full, partial, and total unity
- 2.2 Relations between synchronic and diachronic unity
- 3.1 Bayne's "underlap" model
- 3.2 Illustration of the single state difference
- 3.3 Multitrack vs. single-track conceptions
- 4.1 HOTs and the phenomenal sense of unity
- 4.2 The core idea of the mélange model
- 4.3 A more sophisticated version of the mélange model
- 5.1 Type-K, type-E, and type-H wholes
- 5.2 Two wholes satisfying a relaxed version of globality
- 5.3 Two collections of points
- 5.4 A phenomenal field with centrality-related parts
- 5.5 Atomistic vs. holistic views
- 5.6 A chain of four strongly connected wholes
- 5.7 A combined trajectory
- 5.8 A hierarchy of connected wholes
- 6.1 The point of view defined by experiential contents
- 7.1 Free-energy minimization through perception and action
- 7.2 A counterfactual density
- 9.1 Simple and sophisticated globality
- 9.2 A hierarchy of wholes

List of Boxes

- Box2.1 Definition: Phenomenal Field
- Box2.2 Definition: Phenomenal Part
- Box2.3 The Problem of Phenomenal Unity (PPU)
- Box3.1 The Single State Conception (SSC) of Phenomenal Unity
- Box8.1 The Surprisal Principle

Preface and Acknowledgments

When I started thinking about phenomenal unity as a graduate student, I had the impression that it was one of the most important problems surrounding consciousness, but at the same time it seemed completely elusive. A few years later, when I submitted an early version of this book as a dissertation, I still had the feeling that I could not quite see the path to an adequate solution to the problem, although I had a firmer understanding of what could be said about phenomenal unity in conceptually coherent ways. Since then, the manuscript has undergone a thorough transformation, and I hope the result shows that I have been able to shed significantly more light on some of the puzzles related to phenomenal unity and consciousness.

A lot of people have helped me throughout this process. First of all, I thank my supervisor Thomas Metzinger. He has been my teacher and mentor for many years and encouraged and supported me during all stages of my academic career so far. His comments on chapter drafts were always helpful and at some points enabled me to see problems from a completely different perspective. I also gained a more sophisticated perspective on various aspects of the topic by learning about predictive processing. For this I am extremely grateful to my second supervisor, Anil Seth; his work on predictive processing pointed me to relevant connections between theories of perception and other work in psychology, neuroscience, and philosophy of mind. His rigorous way of integrating philosophical and empirical findings has served as a model for many parts of this book.

I am greatly indebted to the Barbara Wengeler Foundation for supporting me with a PhD scholarship for three years (2012–2015), for awarding me with a Barbara Wengeler Prize for my dissertation in 2015, and for funding trips to several conferences. This has provided ideal working conditions for me, because I have been able to focus entirely on my research. I also profited greatly from the many MIND meetings I attended in the past years. I am grateful to the Barbara Wengeler Foundation for making these regular workshops on current issues of philosophy and related disciplines possible. I am also grateful to Jennifer Windt, who has been the manager of the MIND Group for many years and who organized most of the MIND meetings from which I profited so much.

I thank Philip Laughlin at the MIT Press very much for his support and patience, and his assistant, Chris Eyer, for his support. I am grateful to Amy Hendrickson for creating the wonderful \LaTeX template and documentation at MIT Press. I am also highly grateful to two anonymous referees and Jakub Limanowski for reviewing the manuscript for the MIT Press. Their feedback was extremely constructive and helpful during the revision process.

At Mainz University, I have learned a lot in regular meetings with my “academic half-sisters” Regina Fabry, Ying-Tung Lin, Iuliia Pliushch, and Lisa Quadt. Thank you all for many intensive discussions and conversations. I also thank everyone who participated in the Journal Club and the *Predictive Tiger Club*—those meetings have been helpful and motivating.

Some parts of this book have profited from discussions with colleagues during conferences, and many personal conversations have shaped my thinking while working on this project. I thank (in alphabetical order) David Baßler, Alessio Bucci, Christopher Burr, Joe Dewhurst, Krzysztof Dolega, Adrian Downey, Sascha Fink, Karl Friston, Jakob Hohwy, Max Jones, Marius Jung, Michael Madary, Albert Newen, Oliver Rashbrook-Cooper, Tobias Schlicht, Maria Sekatskaya, Yann Wilhelm, and many others.

My thanks to Ursula Wiese for her valuable assistance in proof-reading. Finally, but most essentially, my heartfelt gratitude goes to Kerstin Kloster-Wiese for being there.

Wanja Wiese
May 2017

I Setting the Scene

1 Introduction

This chapter provides an overview of this book, starting with a summary of the account of phenomenal unity developed here. The summary also contains pointers to all other chapters. Readers who are interested in a particular part of the account will find the relevant reference here. The other sections of the chapter outline and discuss some philosophical background assumptions (e.g., on representationalism, access unity, and interdisciplinary research). Definitions of important concepts can be found in the Glossary near the end of the book.

1.1 A Brief Sketch of the Account Developed and Defended in This Book

This book develops an account of phenomenal unity. It is based on the simple assumption that the concept of *phenomenal unity* has two central semantic components, which I call *phenomenality* and *globality*. The first means that phenomenal unity is experientially manifest; phenomenal unity makes a phenomenal difference in conscious experience (or in “what it is like” to experience things). The second implies that not just particular things we experience are experienced in this way, but *everything* we experience is phenomenally unified.

While these two components are compatible with each other, it is not obvious whether ordinary conscious experience¹ can aptly be described as displaying phenomenal unity in this sense. To clarify and evaluate this claim, I define the problem of phenomenal unity, as well as concepts that are relevant to its definition, in chapter 2. This will allow us to formulate more clearly what is at issue when talking about phenomenal unity. The first problem of phenomenal unity (1PPU) is primarily a phenomenological problem and consists of providing a phenomenological characterization of the phenomenal difference brought about by phenomenal unity. The second problem of phenomenal unity (2PPU) consists of analyzing this phenomenal difference on subpersonal levels of description (such as the representationalist, computational, functional, and neurobiological levels of description).

In chapter 3, I discuss a common (mis)conception of phenomenal unity, which I call the *single state conception* (SSC). According to SSC, what a subject is experiencing at a time is phenomenally unified just in case the subject enjoys a single phenomenal state (which subsumes the subject’s other states). I argue that SSC does not solve the problem of phenomenal unity. In particular, I argue that while it satisfies *globality*, it does not satisfy *phenomenality*. The lesson I draw from this discussion is that these two features, or constraints, as I call them, create a tension when they are applied to accounts of phenomenal unity. Although they do not contradict each other, it is doubtful that phenomenologically plausible accounts of phenomenal unity can satisfy them both.

Some authors argue that accounts of phenomenal unity *cannot* satisfy phenomenality without running into a vicious regress. I discuss this challenge in chapter 4 and will show that there are at least four different phenomenologically plausible ways in which the regress can be avoided; what is more, these options can be at least partially integrated.

Having argued that the phenomenality constraint does not *have* to be relaxed, and keeping in mind that what is at issue is understanding *phenomenal* unity, I next argue that the globality constraint should be relaxed in order to resolve the tension between the two constraints. It can be relaxed in different ways (see chapter 5), and we will see that there is no reason to commit to a single way of relaxing globality. Hence, different accounts of phenomenal unity are possible—in other words, the relation of phenomenal unity is no longer unique: there can be different experientially manifest relations, each of which satisfies a different relaxed version of globality. Moreover, different subglobal phenomenal unity relations can be present at the same time. I call this principle *MULTIPLICITY*. *MULTIPLICITY* implies that phenomenal unity is not a single relation: many phenomenal unity relations exist (the same principle has been proposed by Christopher Hill; see Hill, 1991, 2014).

To do justice to this principle, I suggest that it is more appropriate to call instances of phenomenal unity *experienced wholeness* (at least for many types of phenomenal unity). This reflects one way in which globality can be relaxed, namely, by turning globality into a gradual feature: some things we experience are wholes (like objects or groups of objects), and these can be more or less strong (a perceived group of objects is an experienced whole, but in a weaker sense than the objects it comprises). I call this principle *GRADUALITY*. In general, this suggests that phenomenal unity (understood in this slightly revisionist sense, i.e., by relaxing globality) is more aptly described as experienced wholeness or phenomenal holism.² A third principle associated with this view is what I call *INVARIANCE*. *INVARIANCE* captures the fact that many experienced wholes are constituted by features that are experienced as being invariant. Furthermore, the experienced world and the experienced self are candidates for *globally* invariant experienced parts, because they may be experienced as being connected to everything else a subject is experiencing, and this connectedness may be an invariant feature of many types of experience.

After having explored these different directions in which conceptually coherent accounts of phenomenal unity can be developed, I sketch what I call the *regularity account of phenomenal unity* (RPU). According to this account, a central reason consciousness is described as phenomenally unified is that we experience regularities, which connect experienced features, objects, and events at different spatial and temporal scales. The result is a phenomenal holism, in the sense that we experience multiple wholes at the same time, and many of them are hierarchically nested (an important precursor of this is an idea in Metzinger, 2004 [2003], pp. 143–144).

In chapter 6, I discuss how this and the ideas developed in the previous chapters can be analyzed and enriched on the representationalist level of description. In particular, I consider what benefits a holistic hierarchy of representations can have for a system like the human brain and how formal measures of dynamical complexity (like integrated information and causal density) may serve to operationalize the notion of representational holism.

After that, I combine these ideas with research on predictive processing. Chapter 7 provides a general introduction to the framework of predictive processing and Karl Friston’s free-energy principle, and discusses implications for the relation between action and perception, representationalism, and the disputed (in)directness of perception. In chapter 8, I go on to show how predictive processing can be connected to research on consciousness and attention. In chapter 9 I draw on formal concepts of this framework to provide a rigorous formulation of RPU. I show that at least many experienced wholes can be analyzed with respect to a single computational principle (viz., the tracking of *connecting regularities*). I define connecting regularities as random variables that are predictive of multiple other variables at the same time. The extent to which they are predictive of other variables determines the experienced degree of wholeness. The latter can partially be influenced by attention.

In the final chapter (chapter 10), I recapitulate the most important results of the discussion in this book. In addition, I provide some suggestions for future research.

1.2 Reflection on Methods

As mentioned above, I divide the problem of phenomenal unity into two main problems:

- (1PPU) How can the phenomenal difference (between unified and disunified conscious experiences) be characterized phenomenologically?
- (2PPU) How can the phenomenal difference be analyzed representationally, functionally, computationally, and neurobiologically?

The first question is mainly conceptual and is targeted at the personal level of description; the second is mainly empirical and is targeted at subpersonal levels of description. This is not to deny that empirical results can guide our search for adequate (or more detailed) phenomenological descriptions, and analyses at subpersonal levels are often conceptually loaded. But in principle, the first question can be dealt with from the armchair and mainly aims at conceptual clarity and

logical consistency. Another way of phrasing the question would be: how can we conceive of phenomena unified in a way that is conceptually coherent? Answers to this question consist of descriptions of different *possibilities*. The second question could be expressed as follows: How can these possibilities be realized by a computational system? How can those computations be implemented by a neural system? Which of these possibilities *are* realized in the human brain?

1.2.1 Why an Interdisciplinary Account?

To many readers, it may be obvious that there are interesting questions concerning the problem of phenomenal unity that are not purely conceptual and cannot be answered by armchair considerations. Complementing philosophical accounts by empirically grounded theories may thus be the natural way to proceed. Others might have doubts: Is it really necessary to complement answers to the first question with answers to the second question? Is it even possible and useful to do that? In other words, why should we strive to find an interdisciplinary solution to the problem of phenomenal unity? Why not just stick to pure philosophical reflection?

It should be obvious that it is not necessary to deal with empirical questions. Barry Dainton's (2000) seminal monograph *Stream of Consciousness* is an excellent example of a penetrating, purely philosophical treatment of the problem of the unity of consciousness. Does my approach suggest that something is missing in works like Dainton's? Not necessarily. I do not claim that we have to develop interdisciplinary accounts to solve the problem of phenomenal unity. I just want to suggest that it opens new possibilities. Since researchers are already working on the second problem of phenomenal unity, it will be fruitful to try to connect possible solutions to the first and second problems. A good example is Giulio Tononi's *integrated information theory* (Oizumi et al., 2014).

This theory starts with "phenomenological axioms," which are meant to describe essential aspects of conscious experience. The theory then provides a formal analysis of these axioms. Since the unity of consciousness is one of the theory's explananda, we can ask: Does Tononi's theory provide an adequate informational analysis of key features of the unity of consciousness, as identified by philosophers like Barry Dainton? To answer this and similar questions, it is relevant to bring results from different levels of description together, to see how much we have already learned about the problem of phenomenal unity and to determine which problems and questions are still open as tasks for future research.

1.2.2 How Is an Interdisciplinary Account of Phenomenal Unity Possible?

Some may doubt that this even makes sense: either you provide a phenomenological account of the unity of consciousness, or you analyze, in mathematical terms, how systems like the brain integrate information. These are two completely different projects, and the results will be irrelevant to each other, because the methods and conceptual tools by which they were obtained are paradigmatically different. In other words, skeptics could ask: Are answers to the two questions even commensurable? How can we determine which answers to the second question are compatible with answers to the first one?

The simple answer to these questions is that the proof is in the pudding. The strategy followed in this book is to provide descriptions of phenomena related to phenomenal unity on multiple levels of descriptions, and to take empirical results into account to gain an additional understanding of what aspects of experienced wholes and unity can systematically vary in which ways. This approach is very similar to Thomas Metzinger's (2004 [2003]) method of interdisciplinary constraint satisfaction (as it has been called by Weisberg, 2005). A difference is that I specifically explore the relevance of research in theoretical neuroscience, especially the framework of predictive processing, which is increasingly gaining traction in cognitive science and philosophy.

1.3 Phenomenal Unity and Access Unity

1.3.1 What Aspects of Conscious Experience Are Accessible in the First Place?

A background assumption of this project is that all that is available, from the first-person perspective, are *phenomenal properties*. Phenomenal properties determine "what it's like" (Nagel, 1974) to be in a given conscious state. Since I am a representationalist (see below), I assume that phenomenal properties correspond to the contents of phenomenal representations. More specifically, I assume a form of *weak representationalism* (see Bayne, 2010, p. 52).³ This means that I do not *identify* phenomenal properties with representational contents, but only commit to the weaker claim that any difference in phenomenal properties goes along with a difference in representational contents, and any difference in the contents of phenomenal representations goes along with a difference in phenomenal properties. The content here is *narrow* content (i.e., it only depends on properties of the conscious organism).⁴ Furthermore, the content of a phenomenal representation must be distinguished from its *target* (Cummins, 1996). A work of art that represents an artist as a young man has a target (the person it refers to, viz., the artist) and a content (*young man*). The target of a representation does not necessarily exist (as when I am thinking of a unicorn), but the content always does (cf. Meinong, 1899, p. 187).

The view that only phenomenal properties are accessible from the first-person perspective is not uncontroversial.⁵ For instance, some authors claim that the attentional structure of ordinary conscious experience is also experientially manifest, and that attentional structure cannot be reduced to phenomenal properties (or the contents of consciousness; see Watzl, 2011, 2014). Similarly, some authors distinguish between phenomenal *content* and phenomenal *manner*, which jointly determine phenomenal character, (see Chudnoff, 2013, pp. 565–566). The view that the phenomenal character of a given conscious state is exhausted by phenomenal properties (or the contents of phenomenal representations) cannot obviously account for these apparent, subtle features of consciousness.

However, it is not clear how these putative features of consciousness contribute to the phenomenal character of experience, so it seems that we have to treat them as primitive features. In contrast, a representationalist approach has more resources to clarify—for instance, the experiential contribution of attention—because what we experience is not exhausted by the phenomenal properties corresponding to contents of *first-order* phenomenal representations. In addition, there are also phenomenal *metarepresentations*.⁶ Such meta-representations plausibly underlie our faculty of introspection (see Metzinger, 2004 [2003], pp. 32–37) and many aspects of our experienced first-person perspective, for instance, experiencing the ability to volitionally shift the focus of attention (see section 6.2.4).

1.3.2 Access Unity

According to Bayne's and Chalmers's definition, two contents of consciousness are access unified just in case their conjunction is accessible in the conscious system (Bayne and Chalmers, 2003, pp. 29–32). Furthermore, a mental state "is access-conscious if by virtue of having the state, the content of the state is available for verbal report, for rational inference, and for the deliberate control of behavior" (Bayne and Chalmers, 2003, p. 28). This definition draws on Block's (1995) distinction between phenomenal consciousness and access consciousness. Initially, this is just a conceptual distinction: phenomenal consciousness is defined in terms of "what it's like," access consciousness in terms of (cognitive) functions. However, there is reason to doubt that phenomenally experiencing something without having (some kind of) access to the experienced content is possible. In particular, the seeming richness of phenomenal consciousness, which some authors try to explain by postulating that phenomenal consciousness overflows (cognitive) access (see Block, 2011), can be accounted for without assuming an empirical dissociation between phenomenality and access (see Cohen et al., 2016a,b).

If this is correct, it also means that phenomenal unity always goes along with access unity: when two experienced features, objects, or events are phenomenally unified, the corresponding contents are jointly accessible. In other words, there is a content that somehow corresponds to the conjunction of the individual contents and is accessible (otherwise, no phenomenal difference would be brought about by phenomenal unity). This idea is a central part of my *regularity account of phenomenal unity* (RPU); see chapter 9.

Building the conceptual foundations for this idea, countering possible objections, and providing formal concepts to clarify the idea is the main goal of this book. Another is to provide a survey of research on phenomenal unity as well as conceptual tools and suggestions for how to tackle the problem of phenomenal unity from an interdisciplinary stance.

1.4 Why Anti-Representationalists Should Read This Book

Note that most parts of this book do not hinge on my representationalist background assumption. Core considerations and contributions of this book are orthogonal to any representationalist claims. In particular, chapters 2–5, 7, and most of 8 do not rest on this assumption. The assumption only explicitly enters the discussions in chapters 6 and 9 (and some parts of chapter 8).

In chapters 2–5, I mainly discuss phenomenological and conceptual problems; I also relate the problem of phenomenal unity to attention, the experience of temporal flow, and the subjectivity of consciousness. Moreover, although I spend some pages emphasizing the role of representations in predictive processing

(see chapter 7), you may find it interesting or thought provoking to see how the perspective provided by predictive processing captures attention, action, and perception, which may point to interesting relations between phenomena that are otherwise easily overlooked. As for the “representational gloss” I provide additionally, I hope it increases the coherence of my treatment, but it is not essential to all parts of the overall project.

Notes

1. I am using the terms “conscious experience,” “phenomenal state,” “conscious state,” and “experience” synonymously in this book.
2. The term “holism” has some connotations that are not relevant here, but it avoids the reference to *oneness*, which is suggested by the term “unity.” This is why I prefer this term. However, there are crucial differences to other existing accounts of “phenomenal holism” (see Dainton, 2010; Chudnoff, 2013). These I discuss in chapter 5.
3. This must be distinguished from the (even weaker) notion of weak representationalism defined in, e.g., Chalmers (2004, p. 157), according to which “phenomenal experiences (in a given class) always have representational content.” My own view is that all phenomenal properties can be exhaustively characterized in terms of representational contents.
4. In particular, this also means that I assume local supervenience for phenomenal properties: there cannot be changes in an organism’s phenomenal properties without a change in its physical properties; changes in the organism’s environment are only relevant to phenomenal properties if they go along with changes in the organism’s physical properties.
5. Just to be clear: I am not a sense data theorist.
6. By this, I do not wish to endorse a higher-order approach to consciousness, so I do not claim that a conscious state necessarily involves a metarepresentation. I only claim that what we experience often corresponds to the content of a metarepresentation, but many phenomenal representations are first-order representations.