

Introduction

A generation ago, the problem of intentionality was at the center of activity in the philosophy of mind. Indeed, the project of finding a place for intentionality in the natural order – ‘naturalizing intentionality’ – consumed more intellectual energy than virtually any other issue in philosophy. Yet today the naturalizing intentionality research program generates almost no discernible activity. Why the remarkable turnaround?

There appear to be two conflicting explanations. The first is that the exciting part of the research program has been exhausted, and successfully so, with only technical issues remaining to be resolved. The work done a generation ago outlines the solution to the philosophical core of the problem of intentionality: the reason there is intentionality in nature is that some physical items can track others, with the relevant tracking consisting in a scientifically acceptable relation grounded in familiar, unmysterious causal processes. Some issues remain to be addressed: what the best information-theoretic gloss on the tracking relation is; whether the information-theoretic story should be enriched with a teleological dimension; how exactly to account for intentional failure, i.e., misrepresentation; and so on. But these questions are insufficiently titillating to garner the intellectual energy that the research program had enjoyed a generation ago. Hence the felt drop in levels of activity in pursuit of the program.

The other explanation is less sanguine. It is that the naturalizing intentionality research program bears all the hallmarks of a degenerating research program. The program has run up against principled obstacles it seems unable to surmount. Far from being technical, the problems just mentioned are fatal: every information-theoretic gloss offered thus far has turned out to assign counter-intuitive contents in central, non-artificial cases; it is unclear that teleological elements are genuinely explanatory; and tracking relations simply do not provide the resources to accommodate misrepresentation. Instead of addressing these principled difficulties, proponents of the program have chosen to brand them as technical and proceed to effectively ignore them. As a result, no new

illumination of intentionality has been offered in years, and whatever work does take place has tended to pad the protective belt of the program more than anything else.

Many critics of the naturalizing intentionality research program see a deep reason for its manifest inadequacy, namely, that it ignores entirely the role of phenomenal consciousness and the first-person perspective in its theorizing about intentionality. In doing so, the program has been in lockstep with philosophy of mind at large. A generation ago, it was not uncommon for philosophers to water down the phenomena so that their more naturalistically recalcitrant aspects – notably, phenomenal and subjective aspects – are defined away. However, now that consciousness and the first-person are recognized again as real and central features of the mind, this is no longer tenable. An understanding of intentionality that takes them into account, and grants them the central place they deserve, is well overdue.

On this diagnosis, the naturalizing intentionality program has suffered from a systemic blindspot. It conceives of intentionality as a theoretical posit of sorts, a property ascribed from the third-person perspective in the context of trying to explain and predict the behavior of persons and other intelligibly-behaving systems. The paradigms of intentionality, so conceived, are the propositional attitudes, especially belief and desire. But this ignores a whole other layer of intentionality that is much closer to home and stands at the heart of any natural conception of our cognitive relationship to the world. Once we adopt a first-person perspective, we can see intentionality as a property we sometimes ascribe to conscious states not because it is theoretically profitable to do so, but simply because they present themselves to us in our personal experience as *directed* in the relevant sense. The paradigms of intentionality so conceived are typically not belief and desire, but manifestly intentional conscious states, such as perceptual experiences and conscious occurrent thoughts. An adequate account of intentionality should, according to these critics, be organized around these conscious experiences, not brute and blind tracking relations between brains and their environments.

While research that focuses on the intentionality of phenomenally conscious states has been taking place at the margins of philosophical activity for at least two decades (as well as, effectively, among Austro-Hungarian philosophers toward the end of the nineteenth century), it is only relatively recently that such research has started to

consolidate into a barely discernible competing research program, which we may call the ‘phenomenal intentionality research program.’ The foundational insight of this alternative program is that there is a kind of intentionality proper to phenomenally conscious states that is both distinctive and basic. It is distinctive in that it has key features absent from the intentionality exhibited by any items devoid of phenomenal consciousness. It is basic in that the intentionality of non-conscious items is somehow grounded in, or at least dependent upon, the intentionality of conscious ones. A research program for understanding intentionality that was founded on recognition of the basicness and distinctiveness of the intentionality proper to phenomenally conscious states would be more likely to issue in an accurate and honest portrayal of its subject matter, without letting antecedent ideological commitments pervert the phenomena. And the upshot of that portrayal is that the reason there is intentionality in the world is not that certain mental states bear the right kind of tracking relations to the environment, but that certain mental states exhibit the right kind of phenomenal consciousness, that is, have the right kind of phenomenal character.

What we have here, in fact, are two seemingly opposed outlooks on the source of intentionality. One view is that the source of intentionality is the right kind of tracking relation, the other that it is the right kind of phenomenal character. On the first view, intentionality is originally injected into the world, so to speak, with the appearance in nature of a certain kind of tracking relation. It is when the relevant tracking relation occurs between distinct events or states (including brain states and environmental states) that intentionality makes its first appearance on the scene. Once intentionality has been thus injected into the world, it can start being ‘passed around’ to non-tracking items. But the ultimate source of all intentionality is the relevant kind of tracking relation. On the second view, intentionality is originally injected into the world with the appearance of a certain kind of phenomenal character. It is when the relevant phenomenal character shows up that intentionality makes its first appearance. Here too, once this phenomenal character appears, and brings in its train original intentionality, intentionality can be ‘passed around’ to items lacking any phenomenal character. But the ultimate source of all intentionality is the relevant phenomenal character.

Personally, I am attracted to both outlooks, and so am keen to consider in what ways they may be reconciled or even combined. I am interested in finding ways to place phenomenal consciousness at the heart of the theory of intentionality while at the same time priming intentionality for its eventual naturalization and demystification. When a challenger research program is ascendant, in a first stage it is typically confrontational and rejectionist – it is rebellious. But in a second stage, it is often recognized that research programs rarely have no grain of wisdom to them whatsoever, and that the correct approach is to extract the various grains of wisdom, the fundamentally sound insights behind competing research programs, and integrate them within a single stable overarching framework. I am keen to find a framework that gives pride of place to both consciousness and naturalization.

This book presents and defends a framework that attempts to do so. The framework comprises three central elements. First, it argues that all intentionality is grounded – has its origins – in the intentionality proper to conscious experiences (phenomenally conscious states). Secondly, however, it offers a naturalistic account of the intentionality proper to conscious experiences, indeed an account that appeals centrally to tracking relations. Thirdly, it offers a naturalistic account of the manner by which the intentionality proper to conscious experiences grounds all other intentionality, hence a naturalistic account of non-conscious, non-experiential intentionality. The result is a framework where all intentionality is naturalized, but also where the naturalized non-experiential intentionality is seen to be grounded in a naturalized experiential intentionality.

This framework is developed in Chapters 1, 2, and 4 of the book – with each of these chapters handling one of the three central elements of the framework. Chapter 3 examines an alternative, vaguely non-naturalistic account of experiential intentionality, which has obvious disadvantages relative to, but also certain advantages over, the naturalistic account from Chapter 2. Chapter 5 assembles the results of the previous chapters to formulate an overall statement of the emerging picture of intentionality. Let me close with a fuller roadmap to the book.

In Chapter 1, I argue for what I call the ‘experiential origin thesis.’ This is the thesis that our very conception of intentionality, or aboutness, is grounded in our grasp of

experiential intentionality, the aboutness proper to conscious experiences. The argument proceeds by developing a general model of how we form our conception of certain phenomena, a model which, in combination with the fact that we have a special kind of access to our own experiential-intentional states, casts our access to such states as anchoring our conception of intentionality.

In Chapter 2, I articulate and defend what I call the ‘higher-order tracking theory’ of experiential intentionality. According to the higher-order tracking theory, a mental state’s experiential intentionality consists in its being suitably higher-order tracked to track something. After elucidating the notion of tracking, and arguing against first-order tracking accounts of experiential intentionality, I present the case for a higher-order tracking theory, with special attention to a version of it that I call the ‘naturalized self-tracking account.’

In Chapter 3, I examine the case for a very different approach to experiential intentionality, which I call the ‘adverbial theory.’ According to this, a mental state’s experiential intentionality consists not in any elaborate relation (tracking or otherwise) the state bears to anything, but rather in a certain intrinsic modification of the state’s experiential character. A conscious thought or perception of a triangle has its experiential intentionality, on this view, in virtue of being experienced triangle-wise. I motivate the adverbial theory by arguing initially against tracking-based accounts of experiential intentionality, and ultimately against appealing to any relations in such an account. I conclude, however, with skepticism about the naturalistic prospects of the adverbial theory, and end up avowing a degree of credence in it that, while quite high, is not as high as my credence in the higher-order tracking theory.

In Chapter 4, I turn to the task of accounting for non-experiential intentionality. I consider four possible accounts, and endorse one at the expense of the others. According to the endorsed account, which I call ‘interpretivism’ about non-experiential intentionality, an item’s non-experiential intentionality consists in its disposition to provoke in an ideal interpreter under ideal conditions an interpretive experience, namely, the experience of assigning intentional content to that state. After elucidating interpretivism, I argue for its superiority over the competing accounts.

Finally, in Chapter 5, I collect the results of the previous chapters and synthesize them into an overarching picture of the nature of intentionality. In fact, two pictures emerge: one combines the experiential origin thesis, the higher-order tracking theory of experiential intentionality, and interpretivism about non-experiential intentionality; the other combines the experiential origin thesis, the adverbial theory of experiential intentionality, and interpretivism about non-experiential intentionality. Unsurprisingly, I avow greater confidence in the former.

Harlem, U.S.A.

June 2010