

Situating Mental Depth

Robert W. Clowes 

Lisbon Mind and Reasoning Group (ArgLab), Instituto de Filosofia da Nova,
Faculdade de Ciências Sociais e Humanas, Universidade Nova de Lisboa
robert.clowes@gmail.com

Gloria Andrada 

University of California, Los Angeles
Lisbon Mind and Reasoning Group (ArgLab), Instituto de Filosofia da Nova,
Universidade Nova de Lisboa
gloriandrada@gmail.com

Received 30 January 2021; accepted 11 March 2022; published 15 May 2022.

Abstract

Is the mind flat? Chater (2018) has recently argued that it is and that, contrary to traditional psychology and standard folk image, depth of mind is just an illusory confabulation. In this paper, we argue that while there is a kernel of something correct in Chater's thesis, this does not in itself add up to a critique of mental depth *per se*. We use Chater's ideas as a springboard for creating a new understanding of mental depth which builds upon findings in contemporary cognitive science. First, we rely on the predictive processing framework in order to determine a proposed neural contribution to mental depth, specifically in hierarchical predictive knowledge. Second, drawing from an embodied approach to cognition, we argue that mental depth results from the depth of our embodied skills and the situations in which we are embedded. This allows us to introduce to a new *realist* notion of mental depth, one which can only be explained once we attend to the dense patterns of skillful interaction within a rich artefactual and social environment.

Keywords: Mental Depth; Skills; Skill Refinement; Illusionism; Predictive Processing; Embodied Cognition

1. Depth psychology and its critics

Hilary Mantel's Thomas Cromwell is a capacious literary creation. An inhabitant of Tudor England who can comport himself within any social echelon; Cromwell contains multitudes that are all himself. He is the butcher's boy, the soldier, the hired tough, the accountant, the lawyer, the parliamentarian, the diplomat, and the confidant and advisor to Kings—the architect of the English Reformation. Cromwell is an adept social practitioner and manipulator, always in action, an anti-Hamlet. Hamlet thinks, dithers (and quips), dissects himself but cannot act, and in so doing gives us a Western literary and moral archetype of a certain kind of inwardness of mind. Mantel's Cromwell thinks, reflects, quips, and engages in the minds of others. He regrets but finds resolutions, he shapes events, he balks at simplistic revenge, and relentlessly, he acts some more; always planning, plotting, and shaping the landscape of Tudor England to himself. Cromwell is also archetypal—though an unusual version—of the notion of the “rounded” literary figure that depicts something of the contemporary self-understanding of what constitutes a rich mental life. His depth can be taken as a depiction and echo of the complexity and richness of the life of mind, which we all possess and of which we are so intimately acquainted. He captures something about what we—perhaps at our best—take ourselves to be. We, modern human beings, assume ourselves to be creatures of profound inward *mental depth*.

The idea of mental and psychological depth can be traced to intellectual trends of the late 19th century and perhaps especially to the work of Sigmund Freud. The Freudian idea of a structured ‘dynamic’ cognitive economy involving the interplay of conscious, semi-conscious and unconscious parts has given us foundational elements of much contemporary folk-psychology. It continues to shape both the ‘manifest image’ of what it is to have and be a human mind and also lies behind many ideas in contemporary psychology and cognitive science.¹ This should not come as a surprise. Folk-psychology is in a constant interaction with scientific psychology and when new explanatory notions emerge, such as the notion of the unconscious, we frequently incorporate them in our folk explanations and, consequently, they become part of the tools we use for our individual and collective self-understanding.²

¹ Although Freud is now taken as the main exponent of this idea, other contemporary or near contemporary figures such as William James, Carl Jung, Eugene Bleuler, and Pierre Janet all posited different versions of “depth psychology” where the sources of motivation, ideas, and the self could be traced to the operation of hidden, unconscious, or (sometimes) subconscious forces of which the subject is unaware.

² The notion of the unconscious deeply influenced folk-psychology in the late 19th and early 20th Century both through general cultural discussion but also through works of art from the paintings and films of Surrealists such as Salvador Dali, Rene Magritte, Luis Buñuel, and to the work of vastly popular film-makers such as Alfred Hitchcock. It is difficult to describe much of the cultural life of the first half of the 20th Century without reference to the notion of the unconscious. But the ideas also became part of folk-psychology. Many people who came into contact, even quite indirectly with these cultural trends reinterpreted their own mental life in the light of the new ideas of depth psychology.

Even if much of the psychodynamic framework bequeathed by Freud has fallen from favour, the idea of the deep and unconscious background to many of our psychological processes is still undeniably influential in much cognitivist theorizing about the mind.³ In fact, there is a widely regarded historical view that Freud's idea of the unconscious was somewhat reinterpreted by cognitive psychology (see for instance Power and Brewin, 1991, and Westen, 1996),⁴ where there is a widespread commitment to the view that cognitive processing takes place unconsciously or sub-personally.⁵

Summing up, on a depth conception, our conscious minds are framed as being only the tip of a cognitive iceberg, while the lion's share of our mental life, the sources of dreams, our creativity and eureka moments (but also our hidden desires and biases) take place behind and below the functioning of the conscious mind, in what is often described as "the vast ocean of the unconscious."

This vision of mental depth has come in for a sustained challenge in recent times. For some, the tip of the iceberg metaphor is badly mistaken. The central idea is that mental depth, at least as standardly conceived, may be much more appearance than reality: it is a kind of illusion, that we are only able to maintain through the mind's inbuilt blindness to many of its own gaps and absences.⁶ Our mental life may in fact be thinner, less substantial, and much gappier than supposed by the intellectual giants of the nineteenth century, and literary novelists from Flaubert to Mantel. This view occurs in several places in the contemporary cognitive science literature (Dennett, 1991; Blakemore, 2002), but it is perhaps most clearly articulated in Nick Chater's (2018) book *The Mind is Flat*. According to Chater, the kind of mental depth that Hillary Mantel's Cromwell suggests is a self-flattering but largely fictive depiction of the sorts of beings we are. According to Chater our own mental depth is just as illusory as the apparent depth of fictional creations.⁷

³ The development of psychoanalysis in its various schools, but especially those influenced by the central figure of Freud, was based on the core idea of bringing the unconscious into the light in order promote cognitive change in individuals struggling to understand themselves or change their behaviours. It is worth noting that, though drawing from a different perspective, the concept of depth also played an important role in phenomenology, and more particularly, on the influential work of Merleau-Ponty (1962). Merleau-Ponty explored the relation between depth and agency. He introduced the notion of 'primordial depth' as a basic form of human experience. Depth here is not something internal, but a constitutive element of being an embodied subject that is "involved in the world" (1962, 256–7). In section 4.2, we will come back to this, when we motivate our embodied approach to mental depth.

⁴ Thanks to an anonymous reviewer, for encouraging us to clarify this further.

⁵ For example, the computational model of mind in canonical forms (e.g., Fodor's *Language of Thought hypothesis*) presupposed that the computations that take place in the mind are largely unconscious (see Fodor, 1975).

⁶ For some discussion of whether gaps and absences in the self really imply illusions about self, see Clowes & Gärtner, 2020.

⁷ There are of course many authors who argue that the self is fictional, or a sort of illusion (Dennett, 1992; Hume, 1978 [originally 1739]; Metzinger, 2004). But the claim that mental depth is illusory is a distinct—if related—claim, to the claim that the self is an illusion. The idea of the illusion of mental depth goes

On Chater's account, we emerge as very different sorts of creatures from that imagined by the giants of 19th and early twentieth century psychology and literature. Rather than the sources of our cognitive prowess being hidden away, we are instead ceaseless improvisers. Our brains are always engaged in producing one pattern-completing Gestalt at a time. But, more worryingly, we are also ceaseless confabulators, forever making up fictive reasons and motivations for our actions but blind to the gaps and absences which abound in our conscious mental lives. If Chater is right, the sort of deep inward minds we find in, e.g., Hillary Mantel's Cromwell are doubly fictional.⁸ Much of what we take to be the depth of the human mind is itself an illusion or fiction.

Although we believe there is much to be admired in Chater's critique of mental depth, at least as standardly conceived, in this paper we will present the case against his view. We argue that there is a kernel of something correct in Chater's idea of the just-in-time, improvised character of cognition, yet this does not, in itself, add up to a critique of mental depth *per se*. Instead, we use Nick Chater's ideas as a springboard for creating a new understanding of mental depth.

Our account especially draws upon recent work that identifies the origin of many aspects of the human mind through its dependence upon dense patterns of skillful interaction within a rich artefactual and social environment. More precisely, we argue that the characteristic mental depth of the human mind emerges in the practice of skillful actions. It will take a little work to illustrate what we mean by this idea, so we ask the reader to bear with us as we gradually develop our account here. For now, we observe that the notion of depth can be rebuilt in ways that is scientifically progressive but that at the same time retains some elements that are deeply entangled in the image we have of ourselves as minded creatures. As we will show, our renewed notion of mental depth is not best described as an inner mental depth but a depth that is situated, skillful, and active. It is in our skillful interaction with a rich artefactual and social world that our mental depth unfolds.

Our plan is as follows. We begin (§2) by presenting in more detail the illusionist challenge to the traditional notion of mental depth put forward in Chater (2018). In the following section (§3), we develop a vignette regarding the acquisition and refinement of skilled practices designed to help us illustrate what we claim are the *real* sources of mental depth, which we argue is largely to be accounted for in terms of the depth of skilled situated practices. We examine one concrete form of the acquisition of mental depth in a particular domain through learning to play the cello with the Suzuki method. The next section is devoted to our own alternative approach to mental depth (§4). Essentially, we argue that mental depth is real, and we can find its sources in two places: the depth of hierarchical predictive knowledge (§4.1), and the depth our embodied skills and the situ-

beyond the nature of the self to the claim that the mind is itself a more gappy, less coherent entity with less knowledge about its own nature than we generally suppose.

⁸ Thomas Cromwell is a fictional creation in quite a complex sense. He was also a historical figure in the court of Henry VIII. When referring to Cromwell in the paper we are primarily referring to the fictional creation (Mantel, 2011).

ations in which we are embedded (situated and embodied depth) (§4.2). In the final section (§5) we return to our proposal in order to explain why mental depth in its practiced reality should not be understood as a confabulation.

2. Illusionism and the flat mind

Our account of mental depth is developed against the particular backdrop of *illusionism* (Chater, 2018; Frankish, 2016) and more precisely the form of illusionism developed by Nick Chater in his book *The Mind is Flat*. The cornerstone of Chater's argument is that the inner life as we imagine it, or at least as it is depicted in literary novels, but also in cognitive psychology and much folk psychology is largely illusory. A central target of Chater's book is therefore to debunk a certain conception of *depth psychology*. For Chater, the conscious mind is not the tip of the iceberg with the main edifice of our thought hidden away (see Chater, 2018, p. 186). Rather, for him our brains are Gestaltist parallel processing systems that produce just one coherent thought, perception, and interpretation of the world—as a conscious deliverance—at a time. The sense we might have of a dense background of thought behind this, which only occasionally percolates to the surface, is illusory. Rather our apparently deep thoughts are forged in the moment we encounter a rich and complex world in need of interpretation.

Chater's aim is to show us that the folk picture of mind is wrong. Our privileged access to a private inner world of self is much less firmly grounded than folk psychology regards it as being. Human consciousness is more gappy, improvised, and low bandwidth than we (the folk) suspect. Chater's argument is a form of illusionism, in that he claims that we are mistaken, and operating under an illusion about the nature of our minds and cognitive processes. In a sense, the depth of the human mind is just as much a fiction as the literary creation of Thomas Cromwell.⁹

On the face of it, such a view may seem radically at odds not just with folk-psychology, but also with much phenomenology. Let us look at these ideas and some of reasons for them in a more articulated way.

2.1. The case for the depth illusion and why Chater argues the mind is flat

Chater uses a series of examples to build the case that our sense of mental depth and even the apparent *coherence* of thought are illusory. Two examples here will suffice to make his case clear. One involves our sense of the reality of fictional creations. Chater's book—rather like our paper—begins with the extended discussion of a fictional creation,

⁹ The exact sort of illusion that Chater holds to be the case is close to that originally stated by Dennett in his book *Consciousness Explained* (Dennett, 1991) especially where Dennett's ideas shade into the idea of the Grand Illusion (O'Regan, 2002). As we read Chater, his views are a little more distant from some other contemporary forms of illusionism that hold either that there is no stream of consciousness (Blackmore, 2002), or that qualia is an illusion (Frankish, 2016).

in his case Anna Karenina. Anna is the titular figure of the novel, and one of its central characters, and in the terms, we are using here a *deep* literary figure.¹⁰ Chater points out that although the attentive reader of Tolstoy's novel will certainly feel themselves to have a deep acquaintance and understanding of Anna, it may then come as a surprise to learn that many characteristics—it must be said, rather superficial characteristics of Anna such as her hair color, her height, her build etc.—are never explicitly stated in the book. However, when on reflection we discover that we do not know these characteristics, and indeed that we never noticed that we do not know these characteristics, we are surprised.¹¹ This is an indication that although we think we know how Anna looks and many other aspects of her person, this is an illusion.

The second example draws upon another series of novels, in this case, Gormenghast (see Chater, 2018, p. 21). The Gormenghast novels are famed for, amongst other things, the detail of the descriptions of physical locations and the sense of reality they convey. However, Chater points out that although we may feel we have a deep and coherent sense of the novel's locations and how they fit together, the idea that we can have a fully coherent idea of this cannot be right, for the layout of the castle has been shown to be inconsistent (Chater, 2018, p. 21). The *feeling* of coherence, conveyed by the accumulation of surface detail, makes us think we coherently imagine the castle locale of the novel, but it is an illusion. Crucial for the argument, as it goes for our feelings of the depth and coherence of fictional worlds, so it goes for our feelings of depth and sense of our own inner lives. It may feel like we have a deep coherent inner life behind and below conscious experience but this too—so Chater claims—is illusory.

Chater also develops a series of examples designed to show that our sense of many of the objects of perception and knowledge are much more gappy, sparse, and inconsistent than we typically take them to be. Pride of place in many of these examples is that class of visual illusions that are engendered by impossible figures (see the examples of visual illusions from M.C. Escher in Fig 1.). For Chater, the sorts of experiences engendered by impossible objects rely on an apparent coherence which is not really in the scene, or perhaps better in only local elements of the scene. The global scene is of course, by definition, *incoherent*—at least in the sense of confirming the earthly sense of spatial geometry and gravity. But on closer scrutiny, Chater argues, we find that there is no overall coherence and indeed our sense of depth is itself illusory.

The main idea is that impossible figures are not just mysterious visual games but suggest something profound about our visual systems and the nature of our minds more generally. Our minds tend to project a depth and reality which is not really there. They presage the mind's tendency to mistake the overall coherence and depth of presentations of the world. In this vein, Chater writes:

¹⁰ In nomenclature used by literary theory we could also say she is a “round character.”

¹¹ For a discussion of Dennett's classic paper on surprise as philosophical methodology, see Dennett, 2001.

When viewing an impossible object, we have the overwhelming sense that we are looking at a 3D scene, albeit a peculiar one. But this ‘feeling’ of solidity is completely misguided—we are actually looking at a flat image that has no possible 3D interpretation. This is yet another illustration of the illusions of depth. These illusions of depth, which can be both literal, as with impossible figures, and metaphorical, as with stories and explanations, are everywhere. (Chater, 2018, p. 39).

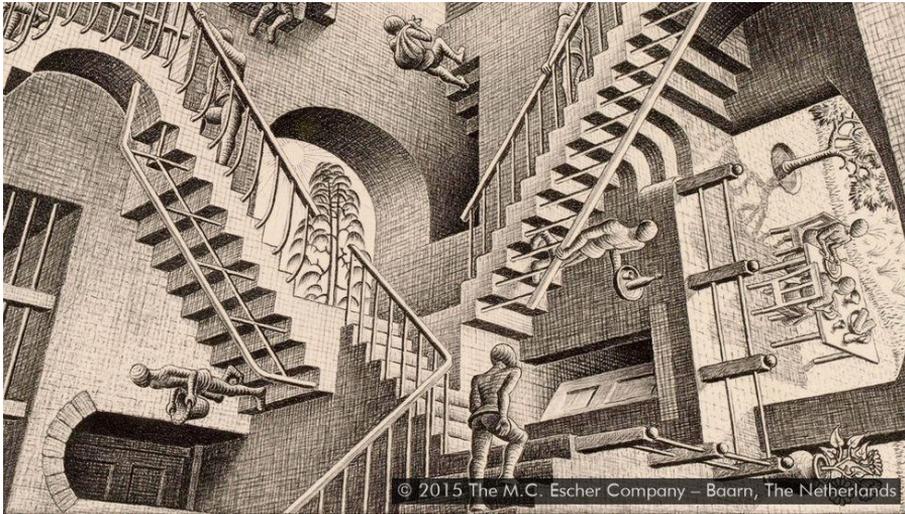


Figure 1

The moral of this illustration is that the sense that we have of the background detail, of proximal depths of the contents of our minds: our perceptions, our thoughts, our memories, and indeed sense of self are all much more like a sense of depth of literary figure like Anna Karenina, or the sense of coherence of an M. C. Escher lithograph. It is apparent but not real. That is why for Chater, the ultimate illusion of depth is the depth of our own mental states, and the depth of our minds.

It is worth noting that these ideas echo previous discussions of gappiness or discontinuities of consciousness such as the idea of *The Grand Illusion* first framed around some puzzling phenomena first discovered in perceptual psychology focusing on phenomena such as change and inattention blindness (e.g., Simons & Levin, 1997; Simons & Rensink, 2005), our experience of the perceptual world is much less high bandwidth and detailed than we take it to be. Marshalling a multitude of such studies, Chater claims that the human mind in general has much less of a grasp on detail than we think, but not just on the detail of perceptual experience but of the content of our minds.¹²

¹² In this respect, Chater concludes: “the mind itself is an impossible object” (p. 21).

According to Chater, the illusion of mental depth is also manifested in the idea of *background processing*. This idea has its antecedents, on the one hand, in the Freudian idea of the unconscious which can be pictured as a background and quasi-personal stream of thought coming to its own inaccessible determinations and holding its own beliefs and desires separate from the conscious access of the ego. On the other, with the idea of background processing which is deeply implied in much traditional (and contemporary) cognitive science. Against this, Chater's mantra is "no background processing," the idea being that the mind is an improviser always making up the best possible, multi-modal and—so far as is possible—integrated interpretation of whatever it is currently being encountered. Instead of there being a constant active background, the brain produces such acts of interpretation one chunk at a time, thereby creating the illusion of a deep mental life.

A core example here is Kekulé's discovery of the structure of Benzene rings, the typical explanation of which Chater takes to be a myth.¹³ After working on the problem of the structure of Benzene rings for many months Kekulé fell asleep gazing into the fire whereupon he had a vision of snakes amid the flames, one of which reached back upon itself and bit its own tail. With this vision Kekulé finds himself wide-awake, inspired and with the solution to the problem of the structure of Benzene alive before his mind's eye. Benzene has, of course, a ring-like / hexagonal structure with each carbon atom bonded to two other carbon atoms and a single hydrogen atom (C₆H₆). Kekulé's discovery of this structure is taken as canonical evidence of the idea of background processing. While Kekulé's conscious mind was resting (indeed asleep) his unconscious mind was said to be working away at the problem. For Chater, this is a sort of post-hoc confabulation of the creative process, and thus the idea of background processing is an illusion.

What then is cognition really like? Chater defends a "cycle of thought" analysis of conscious experience whereby the brain produces one conscious and apparently coherent impression of the sensory deliverances at a time.¹⁴ According to the cycle of thought hypothesis, the mind is able to focus on just one overall interpretation of events and our worldly interactions at a time (we might say a single Gestalt!). The massively parallel connectionist architecture of our brain is constantly tasked with producing one global overall interpretation of the world at a time. All processing power is brought to bear on this, but this leaves nothing left over for independent thought processes going on in the background. A crucial aspect is that the brain is able to handle just one chunk of the world, one perceptual event, one thought at a time. This challenges some central assumptions of cognitivism, which holds that there are lots of processing or "thinking" going on below the level of consciousness.¹⁵ But there is no background processing: just one cycle of thought after another; one element of thought at a time.

¹³ See Chapter 9: The Myth of Unconscious Thought.

¹⁴ For a related early account of the cycle of thought see McCrone (1999).

¹⁵ In connection with how to relate conscious thoughts and the presumed mechanisms that produce them, Chater writes "There are no conscious thoughts and unconscious thoughts; and there are certainly no

The question then becomes: why then do we have this feeling of depth if we do not have it in reality? According to Chater's account, there are two reasons at work here. At one level, our brains are rapid and voluminous improvisors. This means whenever confronted by a particular scenario in need of explanation, our brains are poised and ready to fill in all of the gaps, giving rise to the illusion of mental depth. Much of what we take to be background processing is better explained by a sort of just-in-time filling-in. The other central part of this story is *confabulation*. While our brains are only producing one overall picture or Gestalt at a time, we are at any point able to turn this interpretational process back on ourselves in order to interpret what we must have felt, believed, or thought to achieve the cognitive processes we just did. However, the stories produced in such acts of auto-interpretation or self-explanation, are largely—or perhaps entirely—confabulated.

To summarize, mental life is not some sort of detailed internal picture, or a Cartesian Theatre (1991) as Dennett says. Our mental life is much sketchier and gappier than this (see Chater, 2018, p. 52). In addition, we are largely deceived about the coherence, detail, and even existence of much of our mental life. Chater writes at one point that “the unavoidable conclusion of these findings is that the mind itself is an impossible object” (Chater 2018, 51). Much of the detail of such mental imagery is filled in as and when needed, through more operation of the cycle of thought. Thinking and feeling on this analysis is a sort of improvised, perceptual, just-in-time sort of process. We can only hold one thing in our head at the same time. But we are able to confabulate a back-story for all of our cognitive episodes in line with whatever is our best interpretation and what folk-psychology says. We simply do not notice the many inconsistencies in our mental imagery or even in our visual perception, except where these are pointed out to us where we are shocked or surprised (see also Dennett 2001). We are constantly confabulating a story about our mental processes which is largely at odds with reality. The mind is flat, but we are great confabulators of depth.

thoughts slipping in and out of consciousness. There is just one type of thought, and each such thought has two aspects: a conscious read-out, and unconscious processes generating the read-out.” This formulation is Chater's attempt to circumvent the difficult question of how to relate unconscious processing and what we take to be the conscious contents of the mind at any moment. The problem with this way of expressing things is that it tempts the Dennettian question: who exactly is looking at the read-out? The idea of Cartesian Materialism was developed by Daniel Dennett (1991) precisely to serve as a warning to philosophers and cognitive scientists who assume that there must be a place, conscious experience, where it all comes together. Such a view can become misleading when there is a suggestion that there is an output screen or read out that expresses conscious thought to someone. Who exactly? An inner homunculus? In this paper we do not endorse any particular theory about how content becomes conscious in the brain or to what purpose, but we hold open the Dennettian possibility that there is no place where it all comes together. Perhaps consciousness is, as Dennett suggested, better understood as multiple drafts rather than a single canonical stream. For one way this could be cast into a predictive processing framework, see Dołęga & Dewhurst (2021).

2.2. *The challenge of mental depth.*

Chater's critique lays the groundwork for a novel and challenging view of the mind that seeks to rethink much of what we take to be the nature of the human mind and self. To conclude this section, we want to draw attention to the elements of his view with which we agree and those where we think he goes wrong. We want to read his view, against the grain, ultimately not as an elimination of human mental depth, but as a doorway through which we can see its real sources.

To clarify this, we largely agree with Chater's critique of the Freudian notion of depth and a unified cognitive unconscious. But as we will show, this is not the only notion of mental depth available. In a similar vein, we believe that Chater is largely correct to criticize standard cognitivist views on background processing. However, as will later become clear, mental depth need not be construed as internalist, fully private, or even classically cognitivist in the ways that Chater calls into question. Moreover, his more positive emphasis on the improvisational aspects of mind indicates at least some of the sources of our cognitive prowess. But we do not think that an inference to the lack of mental depth necessarily follows from these insights.

Consequently, our main critique of Chater is not that his picture of a constantly improving mind is wrong, but that the way he understands mental depth undermines the alternative vision he wants to replace it with. In other words, while Chater's idea of the just-in-time, improvised character of cognition is sound, this does not, in itself, add up to a critique of mental depth *per se*. Depth is not the result of serial background non-conscious thinking processes, but its source lies in our rich and deep interpretative capabilities, the depth of the situation and especially the nature and structure of skillful action. Through the development of skills, we are able to attune ourselves to experiencing an expanding set of features of the world, both larger coherent wholes, and episodes of coherence spreading out over longer vistas of time (e.g., Donaldson, 1979, 1992). Essentially, the development of skillful practices reveals new and rich dimensions of the world as we encounter it. Such *mental depth* can be found in the novel and more sophisticated modes of interacting made available through the acquisition and refinement of skills. These practices, moreover, are not illusory nor confabulated. They are real and in the world, and emerge through the intertwined process of coming to rely on the rich affordances of the local environment; especially those made available by artefacts and the patterns of action we build around them.

Building on this, we propose an alternative embodied and situated approach to mental depth. We argue that this approach is immune from the sorts of critique that Chater has staged. Moreover, we stress that mental depth is real, and we can find its sources in two places. First, in our acquired abilities to *see* depth through our accretion of hierarchical predictive knowledge which simultaneously structures how we perceive and act on the world (the perceptual aspect of depth). Second, in the depth of the situations in which we are embedded and to which we learn to skillfully respond. To motivate our account, we will now move to discussing a concrete instance in order to analyze how and where mental depth shows up.

3. Finding depth in skillful practice

In this section, we aim to draw out some of the special aspects of the human mind that show up when it is examined in its ecological setting, namely through the dense patterns of skillful interaction with a rich artefactual and social environment. In particular, we want to illustrate the sorts of skillful situation-engaged minds we have and locate what makes human cognition special by examining the appropriation and refinement of skills. Human cognition takes place amid a set of densely integrated interactions orchestrated between human beings, their artefacts and each other, in a rich cultural environment. Such settings can make possible capabilities of the human mind which appear absent in the laboratory (Donald 2001). Often, we are not very conscious of the contribution that the artefacts make, but, at least within cognitive science circles this has started to change (Clark, 2008; Malafouris, 2013; Norman, 2000). Moreover, these interactions do not take place in a vacuum, but are guided—in ways we will shortly explore—by cultural practices (Hutchins, 1995, 2011; Menary, 2007, 2018).

In what follows, we introduce a vignette aimed to help us explore the situated nature of mind, that is, the interleaved interactions of an artefactual and social context in the development and refinement of cognitive depth in a socially and artefactually constrained context. This (phenomenologically informed) vignette is designed to help the reader understand what we mean by the situated depth of the human mind and form the basis of the way we will then seek to face the illusionist challenge more directly. To illustrate these ideas, we explore the concrete example of a child learning to play the cello guided by the Suzuki method. We will examine some of the details of how a child comes to learn to play the cello in this setting, paying particular attention to both the roles of artefacts and the interpersonal social world, and how they work to support the child's musical development. Our focus is on how the child, within this setting, develops a series of interlocked abilities to control and structure a musical performance, and how these abilities go hand in hand with a deepening sophistication and perceptual sensitivity to the many modalities and possibilities of performance and the world of music.

First, let us discuss the cello and the requirements it imposes. The cello—like all fretless stringed instruments—is a demanding teacher. A central set of skills the child must develop is a practical understanding of the cello itself. Making any kind of music on such an instrument requires not just balancing or holding the instrument in a suitable way but also holding and manipulating the bow.¹⁶ Controlling several interrelated degrees of freedom of movement at the same time is required in order to produce a reasonable sound. Playing even a basic melody on the cello is almost certainly a harder challenge than the more usual beginners instruments of recorder, harmonica, or guitar. Finding and sounding the pitch of any particular note (intonation) is a task requiring considerable perceptual sensitivity and motor dexterity. Making this into a musical sound demands not only the skills of controlling the bow, holding the instrument and finding the right notes to play, but also the development of a more sensitive musical ear. That is,

¹⁶ Holding the bow and actually bowing is a task that may take years to accomplish.

perceptual development and growth are required alongside considerable control of movement and posture.

One characteristic way that the Suzuki method seeks to manage the difficulty and complexity of producing a note with good intonation is by placing lines of colored tape (virtual frets!) on the fingerboard of the instrument (fig. 2). These are, in the first instance, placed where the major second, third and perhaps fourth degrees of a major scale would be—considering the open string as the root note—and serve as visual guides to where the student should place their fingers in order to produce good intonation. In this way, the child can make use of the visual cues (tape), so to orient their growing abilities to find the desired pitch of a note through a proprioceptive familiarity of how one's fingers reach for that note in combination with what their musical ear seeks to hear. The tape then serves as a 'scaffold' from which more refined perceptual and active possibilities are developed. Eventually, the visual clues become superseded by the child's proprioceptive knowledge of where to put their fingers alongside the growing ability to *hear* whether a given note is in tune. At this point, the colored tapes are removed from the instrument. This takes some time!



Figure 2

However, the ability to play with good intonation is, if anything, an easier and more natural task than learning to use the bow in such a way as to produce a consistent and well-defined musical sound on a string. This requires, in the first instance, significant attention to posture and movement in order to develop the habits and inculcate a set of skills whereby the child becomes progressively able to control the fluid and seemingly effortless production of musical sound. Learning to do this effectively requires a significant development of motor-control and postural sensitivity on the part of the child. The enveloping socio-cultural setting, at least with the Suzuki method, plays a central and explicit part in developing these abilities.

The Suzuki method depends on the cultivation of a special nurturing relationship between teacher, child, and parent (or caregiver).¹⁷ The parent is introduced to the Suzuki method and theory, and is often encouraged to attend lessons. A guiding idea of this relationship is to bring the child, and importantly the parent, to a gradually deepening sensitivity to musical practice and to musical performance. Part of the reason is to develop a supportive and knowledgeable background to aid the growth of skilled practice and refined perception in the child. The cultivation of this background helps the parent to support the child's home practice time which is crucial to the growth of their abilities to gradually refine their skills more autonomously.

Implicitly, there appears to be something of a necessary distribution of cognitive labor between child and parent.¹⁸ While the (young) child is deeply engaged in the business of holding the bow just so (e.g., respecting the balance point), or resting the cello on a particular point on their chest and left knee, the parent is often busy taking notes to help prompt later practice sessions at home.¹⁹ It is not unusual for the Suzuki student to begin their learning at four or five years of age, so it is certainly the case that the parent needs to take substantial responsibility for explicitly remembering what the child is supposed to be doing, while the child is concerned more with the physical requirements of the cello. As the child's sophistication grows, they are able to take over more of the explicit memory burden making (mental) notes from a weekly lesson in order to control the structure of home practice for themselves, and indeed taking more responsibility for their own targets and goals.

Although the Suzuki method emphasizes how playing music and the production of tone precedes any deep musical theory, at least as presented to the child, as the child learns to inhabit the physical constraints that allow the production of musical sound, they also have to develop their abilities to understand and to hear music. These aspects are crucial and eventually come to be accompanied by substantial theoretical knowledge. Cello lessons are therefore often accompanied with music reading and eventually theory classes, which introduce children to musical notation—alongside many central musical concepts about how to modulate and control performance. Another factor of importance is the “class conjunto”: a joint class where the students play in a mini-cello orchestra and where

¹⁷ A note on the method cited from the cover of a current Cello School book reads that “The Suzuki Method involves the student, the trained teacher, and the parent. Parents work with teachers to create a fun, nurturing environment for learning by attending lessons with their child, serving as “home teachers,” and playing music at home.” (Suzuki, 1991, back cover)

¹⁸ This could be thought of as a case of distributed skill acquisition. See Sutton et al., 2020 for an account on distributed memory and socially distributed remembering.

¹⁹ This can also be looked at as a sort of social division of memory. Whereas the child is (implicitly) learning how it feels to hold the cello and the bow in a way the musical occasion demands—often a process of implicit practical understanding of the posture of the body—the parent may be taking notes of key points, such as what should be remembered or picked out for detailed attention in the practice sessions at home between the weekly sessions with the cello teacher.

children begin to learn about the difficulties and joys of playing with others and also working towards performance in front of an audience.²⁰

Playing the cello and certainly anything approaching mastery requires a great depth of skilled performance. The knowledge required is not just of the theoretical type we have just mentioned, but prior to this, and also grounding it, a practical knowledge of skillful practice involving an attunement to the requirements of the artefact mediated production of sound. The acquisition of the requisite skills is—as we have just described—heavily reliant upon the production of custom environments, a densely socially scaffolded process of acquiring specific skills and, not least, the gradual physical attunement to performance, e.g., adaptation of muscles, sense of bodily posture and even how to manage emotions through performance. These skills are very aptly described as being situated-embodied practices. They require the child to build a sensitivity to certain tools, the cello, the bow, the resin, and how they interact, as well as a deep appreciation of the possibilities and constraints of movement of their own bodies with respect to these implements. Playing well will also require significant muscle control and an appreciation of posture to the point that it becomes a sort of transparent habit. It will eventually also require a working understanding of related and sometimes less practical domains such as more theoretical knowledge (e.g., scales, arpeggios, the cycle of fifths, etc.) and other forms of practical understanding, (such as how to read music, relate to others in the orchestra, play at an appropriate volume, etc.), and how all of these too relate to the production of sound (e.g., the production of vibrato, using the bow).

A dense set of practical skills, theoretical knowledge and situated-embodied practices all have to come together to allow a good performance. In performance, it is gradually no longer the production of an individual note, or successful body posture that is uppermost in a student's mind. These foundational features need to be in part taken for granted and are handled by the habits of skilled practice. It is the performance of quality that is the ultimate goal and eventually many of these hard-earned skills will become, and need to become, transparent in action, so that the student can think of the performance in more abstract terms. The student instead learns to think of musical phrases, or passages of play, or how a whole performance can be subtly modulated and approached in a number of ways, about which the performer can make conscious choices in order to evoke a different emotional response. Indeed, the development of such practices are intimately related to a growing sophistication and perceptual refinement on the part of the student. This perceptual growth can be seen as moving outwards from the production of a single note to an expansion across a number of temporal and spatial scales such that the child's mind is able to focus on more refined, more temporally extended, more abstract and more emotionally affective aspects of performance.²¹

²⁰ In fact, at least in the music school of one of us, children are encouraged to make recital performances for groups of parents and family more or less from the beginning of their musical journey.

²¹ Although we do not have the space to really develop these ideas here, it is possible to cast the child's developing skills in a more general temporal framework, proceeding outwards from what developmental

We have introduced this vignette at length to focus on a particular area of human activity which we think is neglected when we consider mental depth. We believe that it is largely these sorts of practices in which human mental depth inheres. It is essentially a set of skilled and situated practices, dependent on a particular artefactual culture and a variety of social supports, practices and interpersonal relationships. The type of mental depth that is produced can be highly resilient but is also highly situated and depended on particular environmental supports and extended processes of education and enculturation. It is not however in any useful sense illusory or confabulated. In what follows, we will be taking this sort of skilled practices as an archetype of human mental depth.

4. The sources of depth of the skillful mind

We will now demonstrate how it is possible to theorize the acquisition and refinement of skilled mental depth in terms of some particular reference points in contemporary cognitive science.

4.1. Mental depth and the rich interpretation of the situation

Skills, while being interactive and dependent upon environmental props, also require something that the agent brings to the situation. A skilled agent may depend on proximal tools and what these tools afford, but it is undeniable that agents also bring ways of acting in and seeing to that situation. These ways of seeing and acting, these abilities to pick out the unique affordances that only a skillful agent can see are in need of explanation. Our account will make sense of this ability to see and develop sensitivities to new affordances by giving an account of the neuronal contribution to the experience and ongoing activity of mental depth, especially as it appears in the cultivation and exercise of human skillful action. We will first present an account of skill growth drawing from the predictive processing framework (Clark, 2016), before returning to our account of the embodied and situated nature of skills.

Let us begin by reviewing some Predictive Processing basics. Two core elements of predictive processing are especially relevant for our purpose. The first is the *generative model*, whose main task is the prediction of sensory signals, and the second is the *precision-estimation mechanism*. The generative model is a unified body of acquired knowledge based on previous experiences. A central thesis of the predictive processing account is that the way perception works is by predicting bottom-up sensory cues drawn from its best models of what is likely to be causing them. Perception is an active process, and the brain contributes to this activity. The prediction of sensory input drawing from the statistically salient history of the agent is a risky process, that is, it could easily go wrong. That is why there is a second feature in the predictive processing framework that

psychologist Margaret Donaldson call the “point mode” of the here and now to ever expanding vistas of time in which a performance is located (Donaldson, 1979, 1992)

works with the generative model, namely the precision-estimation mechanism. The central idea is that, besides the prediction task, the brain assigns a probability to the source of information given its estimated degree of certainty or uncertainty. The result of the interplay between these two mechanisms is both an optimization of top-down predictions, or *priors*, together with an optimization of the precision-estimations that determine the probability of bottom-up prediction errors in the processing hierarchy. In other words, mismatches between expectation and input (i.e., prediction errors) are propagated 'forward' in the system where they serve to further refine the top-down predictions. Supporting this *top-down* or *knowledge-driven* prediction, we can attend for instance, to the "hollow face illusion" Clark (2015, p. 3767). This illusion is produced because we are used to convex faces in our daily experiences, that is, we *know* faces are convex. According to the predictive processing account of perception, these experiences are accumulated, contributing to our *generative model* of statistically salient experiences. The illusion happens when we are presented with a screen with a concave mask. One way of explaining this effect is by assuming that the expectation of convexness is so strong that we cannot help but experience certain face-like input as convex. This means that we expect faces to be convex, or in more accurate terms, we have a "deep sub-personal 'expectation' of convexness" (Clark, 2015, p. 3767). In what follows, we use these elements (i.e., the generative model and precision estimation) which are shaped by our prior actions and histories of interactions through the development of skillful practice, in order to explain their contributions to mental depth.

Once we reject an intracranialist approach to cognition (more on this in the next section), the brain emerges as part of the larger ensemble that is the embodied cognitive system. This system is in constant interaction with, and dependent upon its changing environment through a series of looping interactions. According to our preferred view on predictive processing, the predictive brain is an extra element of situated embodiment and not a replacement for it (cf Hohwy, 2013). One way of describing simply what happens when we perceive something according to predictive processing, is that "we see the world by (if you will) guessing the world, using the sensory signal to refine and nuance the guessing as we go along" (Clark, 2016, 43). As we have just seen, the process of trying to guess or accurately predict worldly states is accomplished sub-personally by the rich mechanisms afforded by the predictive brain. However, we also act on the world in order to optimize precision-estimation, or in other words, to render it unsurprising. This is achieved by means of what has been called *active inference*.²² As Fabry (2018) writes, in active inference: "embodied actions bring about changes in the available sensory input so as to confirm the accuracy and adequacy of top-down predictions. On this construal, any type of bodily movement—from ocular-motor adjustments to locomotion—has the potential to confirm the best probabilistically generated predictions" (p. 2486). Notice that this role of embodied action in perception is in tune with what we saw in our pre-

²² Hohwy (2013) distinguishes between perceptual inference (i.e., accurately predicting worldly states) and active inference (i.e. acting on the world in order to optimize the precision estimation task). As he notes, both of them are computationally similar, since they are coordinated by the optimization of precision estimations. However, he claims that they have a "different direction of fit" (p. 178).

vious discussion of learning to play the cello and especially how this relates to a growth in perceptual and related conceptual abilities. We will shortly develop this idea further by relating it to an account of how perception is constituted by a practical knowledge of sensorimotor contingencies, or the structured nature of worldly affordances (Noë, 2002).

The picture we offer is the following. The embodied agent generates internal dynamics based on their prior history of interactions mainly in virtue of neural networks that compose the generative model;²³ these endow us with sub-personal mechanisms that infer the likelihood or uncertainty of sensory regularities based on what has been previously encountered. Importantly, we act on the world in order to adjust sensory information with prior predictions. Perception is thus constrained and contextualized by prior action and the history of the agent. This constraint and contextualization forms the neural basis of our experience of mental depth, by allowing us to *recognize* complex and affordances for action in the world, especially in situations in which we have developed skillful mastery.

To illustrate this, let us return to our cello example. Practicing the cello requires the slow development of a series of sensitivities toward the instrument including how to hold the instrument appropriately, how to use the bow, and how to find a note. For a more skilled practitioner some of these aspects become increasingly facile and transparent and the performer can concentrate on other dimensions of performance. If we apply the concepts that the predictive processing framework offers us, we can explain what is at stake in the following way.

Generative models or “the generative model” are the proposed neural realization of what we bring to the situation in order to appropriately respond to it.²⁴ Through the tacit practical knowledge made available through a much-refined generative model, the adept

²³ It has to be admitted here that despite the widespread view that the idea of generative model can be grounded in the structure of the brain, the precise neural implementation of this model, and indeed whether real brains are predictive processing systems remains controversial. Our discussion then of predictive processing in neural terms is of a still uncertain empirical hypothesis, and one that is the subject of ongoing disputation and the search for empirical evidence.

²⁴ We refer to generative models here but note that the term generative model is the more usual and canonical terminology. The terms generative model implies that the whole brain can be seen as one multi-level (hierarchical) and multi-scalar interrelated model of the causal structure of the world with which an organism interacts. Karl Friston’s use of the term generative model also makes a connection to the free energy principle, which sees predictive processing as fundamentally being a way of reducing free energy (See Wiese & Metzinger, 2017 for a primer on free energy principle and how this interacts with hierarchical models). Another idea—and more important to the context of our discussion—is that the generative model with the highest posterior probability, or, put another way, the “winning hypothesis” of the brain about the current causal structure of the world can be used to explain the content of consciousness at that moment (e.g., Hohwy & Seth, 2020). We take both of these important, and possibly revolutionary ideas to be interestingly controversial but largely beyond the scope of this paper. The use of the term ‘generative models’ indicates then that one could hold a predictive processing account of mind without subscribing to either of the two just mentioned propositions with the conceptual implication that it is

cello player sees possibilities for controlled action that are not available to the novice. Put another way, the adept cello player does not see, hold, or hear the same instrument as the novice. Hours of practice and training are reflected in the models that the adept brings to the situation and thus, the affordances that they perceive. The process of practicing can be understood as a process of refinement of generative models. Precision-estimation mechanisms contribute to the refinement that takes place during the process of practice. The propagation of errors and constant attunement contributes to developing and fixating more accurate and refined strategies for the controlled expression of music. Importantly, these processes take place in a sub-personal fashion, giving rise to a form of 'self-supervised learning' (Clark, 2016, p. 18). The idea is simple: by constantly trying to infer the sensory signal via the generative model, our history walks with us in a subtle way, gradually shaping our experience of the world and the depth that we find there. It is the deep history of our skillful interactions—as we achieve mastery of the sensorimotor dependencies in any given domain or type of situation—that determine how we see the affordances therein. The precision-estimation mechanisms contribute to the process of learning and achieving the mastery of skills, by contributing to the accuracy of our skilled gestures in action.

One other aspect of the importance of generative models here is *Active Inference*. Active inference emphasizes how probabilistic beliefs are not only changed by optimization of the generative model, but also through acting *on the world*. One way of controlling the variations of sound produced by bowing is by subtly changing one's grip on the bow. Performing an action in one way or other changes incoming sensory data, and thus allows us to refine different aspects of a performance. A central way in which the predictive processing idea meets up with skilled practice is in how it allows us to give an account of the integral role of movement and action in allowing one's models to produce the optimal grip upon the world in order to produce skilled performance.

We would like to note again here how developing a skill such a cello performance requires a significant growth in perception; a greater sensitivity to the richness and multiple modalities of the world and how we can act upon it (Donaldson, 1979, 1992). One way to think about this is in regard to the two other ways in which predictive processing theory seeks to explain our cognitive abilities. It is said that perception realized through generative models is "richly world revealing" (Clark, 2012). A central idea of predictive processing is that a generative model is a model of causal structures in the world. We do not perceive surfaces as such but the interplay of causal forces that "explain" what our sensory systems receive. Such a view allows us to naturally locate the development of

possible to hold a form of predictive processing view without assuming that all "generative models" in the brain need be integrated into one model. The notion that there may be several predictive processing models in the brain may also comport well with some other theories of consciousness such as Dennett's (1991) multiple drafts theory of consciousness and we draw the reader's attention to a recent paper that assays just this possibility (Dołęga & Dewhurst, 2021). Indeed, since predictive processing currently appears consistent with a number of different theories of consciousness this could be seen as a limitation of the view, at least as a means of explaining consciousness (Schlicht & Dołęga, 2021). Thanks to an anonymous reviewer for encouraging us to clarify this further.

perceptual sophistication and how it is tied to the development of skillful action. It is not just that we develop more sophisticated skillful practices, but these practices reveal to us a growing sense of the intricacy and interplay of forces in the world beyond us. Moreover, predictive processing proposes a much more unified account of perception and action than classical cognitive accounts. These two factors can be very naturally applied to music where it is through the development of skillful practice that much of the depth of music and musical structure is revealed. This close intertwining of action and perception allows us to explain how musical sensitivity grows. It seems natural to link this to the growth of generative models and to use these to account for the expanding richness and temporal range of our perceptual abilities. The growth of sensitivity in the performer can be explained in part through the growth of perceptual depth and its tight reliance upon situated and embodied skillful practices. The predictive processing account puts on a firm foundation what the brain brings to the acquisition of deep skills.

With this in mind, we can go back to Chater's account of mental depth. If our account is on the right track, this means that the claim that there is no background processing is not strictly true. Our claim is that there *is* indeed a form of background processing, which can be accounted for in terms of the hierarchical predictive knowledge with which we perceive and act on the world, and that contributes to eliciting the experience of mental depth. In this respect, depth is partly explained by our prior associations and experience. However, this background processing should not be conceived of in a classically cognitivist and internalist way. As will become clearer in the next section, we understand predictive processing in embodied and situated terms, where the brain is part of a wider embodied system.²⁵

4.2. Situated and embodied depth

Having put this predictive processing account of cognitive architecture in place, we now foreground how mental depth is necessarily situated and embodied and substantially relies upon, and indeed is constituted by our skills. Mental Depth is not something best understood as purely internal, or that happens outside of the conscious mind nor is it a narratively generated confabulation. Our claim, moreover, is that mental depth is not simply the result of neural activity. Rather it emerges through the development of specific practices scaffolded by the social world and through deep enculturation. The development of skilled interactions and artefactual manipulations allow us to make use of the potentialities of the rich material environments in which we are embedded. These and that actively contribute to the feeling of depth and coherence. In the previous section (4.1), we focused on the neural contribution to mental depth in terms of predictive processing. Here we will explain its situated and embodied character and provide the theoretical framework that allows us to ground our account of mental depth in skillful action.

²⁵ See, for instance, Clark, 2015, 2016; Fabry, 2018; Gallagher & Allen, 2018, to mention some. Though the level of embodiment varies, all of them reject the view where predictive processing is simply a matter of the working brain.

Let us begin by observing that the folk image of our psychological lives is that the mind (i.e., one's thoughts, desires, memories, etc.) is located somewhere in the head. This idea receives further support from much cognitive neuroscience, according to which mental and cognitive processes are implemented *only* by the brain and the central nervous system. In this way, according to this folk position, mental processes take place somewhere *in* the brain and the central nervous system. We may call this view *cognitive intracranialism*, and we can locate Chater's account within the space of intracranialist views. Recent approaches in cognitive science challenge cognitive intracranialism by expanding the cognitive realm so as to include the agent's body and the skillful interactions in an environment (Newen, De Bruin, & Gallagher, 2018).

Our account of mental depth rejects *cognitive intracranialism* and instead conceives of cognition, skillful action and mental depth as all strongly embodied and embedded. By strongly embodied, we mean that cognitive processes involve the body acting in and on the environment (Clark, 1997; Gallagher, 2005). By strongly embedded, we understand that some mental and cognitive processes are the result of the integration with states and processes found in the environment (Menary, 2007). We conceive of the environment not simply as the physical environment but also a structured social, artefactual, and cultural environment. This is captured by our reference to the *situatedness* of cognition.

In order to present how our situated approach to mental depth contrasts with Chater's view, we will review Noë's (2002) reply to an illusionist challenge faced by the standard cognitivist approach to (especially) visual perception (Blackmore, 2002; Dennett, 1991). The challenge is the following. We know that the information received by the visual system is fragmented and discontinuous, e.g., we are susceptible to change blindness type illusions. Yet, our visual experience of the world, in a phenomenological sense, generally appears coherent and continuous. However, it is believed that we have "a richly detailed picture-like experience of the world, one that represents the world in sharp focus, uniform detail, and high resolution from the center out to the periphery" (Noë, 2002, p. 2). To explain how that could be the case, one approach is to argue that the brain needs to 'fill in' the missing sensory details. However, there is no definite evidence that the brain actually 'fills in' all the gaps. (Dennett 1991, pp. 344–356). The consequence is that despite the *apparent* continuity of consciousness in experience, consciousness might in fact be discontinuous. Similar to Chater's criticism regarding mental depth, there is thus an asymmetry between common beliefs concerning perceptual experience and what perceptual experience really is.

One way of making sense of this situation is by claiming that we normally fall prey to a continuity illusion (Dennett, 1991). We are victims of an illusion concerning the continuity of experience or the character of visual consciousness. We are thus wrong when it comes to what we take our experiences to be. Resisting this option, a strongly embodied or enacted view on cognition opens up a different possibility (Noë, 2002). The main idea is that the illusionist challenge is built on an incorrect view on what the phenomenology of perceptual experience is.

A central thesis of the embodied approach to cognition is the continuity of perception and action (Clark, 1997; Hurley, 1998; Noë, 2002). This thesis argues against the standard cognitivist approach according to which perception is a passive process, and sensory stimuli are received and processed in order to give rise to a behavioral output. A strongly embodied or enactive approach to cognition, conceives perception as intimately linked to the action of an organism in its environment. One way of articulating this continuity is by considering that perceptual experience is in fact constituted by sensorimotor expectations acquired and developed through the exploration of the environment (Hurley, 1998; Noë, 2004). This rejects the standard construal which, borrowing from Hurley (1998), is committed to the “classical sandwich model”, that is, the claim that cognition is like the filling of a sandwich which is contained by perception and action. Instead, on the strongly embodied construal of cognition, perception and action are intimately connected to the extent that perceiving *is* a form of acting.

These points call into question what the illusionists took perceptual phenomenology to be: we do not perceptually experience the world as a whole in a “snapshot” but we experience having access to the world in virtue of our sensorimotor abilities (Noë, 2004). In other words, perceptual experience is not best understood as a rich perceptual world which is presented or given to our sensory channels, but as a skilled cognitive engagement. Take, for instance, the experience of visually perceiving a wall. We don’t experience the whole of the wall’s surface but we “experience the wall as present and you experience yourself as having access to the wall by looking here, or there, attending here, or there” (Noë 2004, 4). This offers a solution against illusionism: consciousness is not continuous in the standard “snapshot” sense, but this does not entail that it is discontinuous, and thus we are constantly prey to an illusion. What this shows is that we need to move beyond the standard construal of perceptual experience.²⁶

Our strategy against Chater’s challenge to mental depth follows a similar dialectical structure. We resist the claim that mental depth is an illusion, but our resistance only succeeds if we adopt a different and more plausible phenomenology of mental depth.

To make our case, let us briefly recall Chater’s challenge. Chater challenges the folk/standard conception of mental depth, according to which our inner mental life is as deep as it is depicted in literary novels and in the assumptions of folk-psychology. The challenge is grounded in the fact that according to Chater’s view there is no background processing—a claim we have just challenged—and our mental life is in fact sketchy, gappy, and highly incoherent. As we have seen, Chater claims that much of the detail of our mental imagery is filled in as and when needed in a manner that Chater describes as “the cycle of thought.” The consequence that Chater draws from this is that our experience of mental depth is thus illusory. But is this the only option? We don’t think so. In fact,

²⁶ This also allows the explanation of the *perceptual presence* of objects or scenes despite them not being entirely modally present. The idea is the following: when you see a bottle in front of you, you perceive the bottle as being entirely present, despite the fact that from where you stand, you cannot see the back of the bottle. For more on this see Noë, 2002, p. 9.

once we endorse a strongly embodied and embedded approach to cognition, we can offer a new account of mental depth and its phenomenology.

As we have stated earlier in this section, our account of mental depth begins with the consideration that cognition is largely a form of situated skillful activity. In other words, cognition is largely a matter of the embodied manipulation of external structures (Rowlands, 1999). This means that it is in virtue of learning such manipulative skills and the acquisition of social practices that our cognitive abilities are transformed and extended (Menary, 2007, 2012, 2018). Cognition is thus not solely an intracranial activity, but it includes different patterns of skillful action. Understanding cognition requires us to pay attention to the different “skilled gestures” (Clowes, 2019; Andrada, 2020), that is the skilled movements and manipulations that allow us to act and interact with the material culture surrounding us as well as with others²⁷.

Taking seriously such a strongly embodied and situated approach calls for a reconfiguration of our notion of mental depth. It requires us to see that mental depth is not an entirely inner matter, constituted by hidden causes and background processing, but depends upon the kind of access we have to our environment. Especially insofar as our access involves skilled manipulations and sensorimotor interactions. In other words, depth phenomenology is not simply the result of neural activity, but the way the neural activity is embedded in a sensorimotor dynamic involving conceptual growth and environmental exploration.²⁸ This forces us to reconsider what Chater takes the default position concerning the phenomenology of mental depth to be. It is not that we experience ourselves as deep characters with a rich inner and largely hidden mental life, but we discover ourselves as deep in virtue of our situated activities in our active engagement with and responsiveness to the environment. In other words, we are not under the illusion of depth, but depth is not as folk psychology traditionally took it to be.

The phenomenology of depth begins by noting that it is through our activities, exploration, and developing sensitivities to worldly events and occurrences that we discover depth. We find depth in the course of our worldly activities through the sense of multiple ways of acting upon and thus thinking about an object, tool, or situation. We create and experience depth especially in the exercise and mastery of skillful practices, or sometimes through the observation of those practices being performed by others. We find

²⁷ Extended cognition so understood is thus “a theory of *access*” (Krueger, 2014, p. 5). This approach to extended cognition is less ontologically committed and thus less controversial given that cognition is extended, that is, cognition includes extra-organismic elements only when engaged in sensorimotor dynamics.

²⁸ It should be noted that this situated and embodied aspect of mental coherence connects with Merleau-Ponty’s notion of ‘primordial depth’ which is not the result of cognitive calculation but something that is gained and achieved by an agent in virtue of being situated in the world. Depth is thus experienced by a “subject involved in the world” (1962, pp. 256–7). Merleau-Ponty writes that depth “is, so to speak the most ‘existential’ of all dimensions, because... it quite clearly belongs to the perspective and not to things... It announces a certain indissoluble link between things and myself by which I am placed in front of them” (1967, p. 256; quoted in Bredlau, 2010, p. 413).

depth in worldly interactions and practices. Just as our discussion of learning to play the cello suggested, mental depth comes from our ability to attune to and/or respond to these sort of rich practices and structures as we find them in the world.

One way of noticing how our skillful practices in the local environments contribute to the phenomenology of mental depth is by noting what happens when we are deprived of access to some aspect of the enveloping material and social environment upon which we usually depend. In such circumstances we can experience the loss as a sort of deprivation of parts of ourself, or, put another way, as a loss of mental depth (more in the next section). That is why we consider the kind of access we gain in virtue of our situated skillful practices as in fact at least partly constitutive of the experience of mental depth. This access, though it can be, needs not be conscious, and for the most part, will not be the focus of our attention and conscious thought, at least to those practicing the necessary skillful activities. However, there will be a particular phenomenology associated with this constant access which is afforded by our embodied manipulations.

Moreover, conceiving of mental depth as situated and embodied can also help us understand mental coherence in a way that does not require the sort of inwardness of mental life that Chater challenges. Take for instance what Reed's (1993) writes recalling his experience with his patients suffering from traumatic closed head injuries with severe brain damage. Despite their failure to organize coherent skilled behavior (e.g., preparing coffee or pouring water), they were able to preserve certain coherence through their actions (e.g., they mistook liquids but they did not pour liquids on a plate, nor they did put solids in a glass). This can be explained by the fact that they preserved certain basic affordances learned in virtue of their previous interactions within their local environments, and their local environments partly contributed to their mental coherence.

Summing up, the depth of human mental life is based in our skillful practices in the world of the sort we have just assayed in our discussion of the acquisition of skill of playing the cello. Although the acquisition and maintenance of such a skill is in many respects particular and *sui generis*, skills have a sort of general structure which can be used to articulate much of what is specific to the human mind and our particular form of mental depth. Depth is acquired and sustained through the development of skills that involve the elaboration of specific abilities which are often or usually highly dependent upon the tools and artefacts of the human-made environment. We thus are mentally deep not because of inscrutable hidden processes, but because of the way we are able to respond to the world. This situated skillful depth is not illusory, and neither is it in any simple sense inner. Rather, in order to understand both how it arises, and what is its essential nature, we need to understand how it intrinsically depends upon the situated nature of our being in the world.

We thus conclude that depth is not an illusion that results from our conscious or unconscious narratives and confabulation, but it emerges and is present in our skillful practice. Our experience of depth on this model is deeply related to our ability to produce complex, structured, and refined activities. These abilities are rooted in the history of our development of skills and largely resides in the prior associations, expectations (genera-

tive models), and experiences which are built-into the architecture of our perceptual and cognitive embodied mechanisms. The conclusion of this section is that the deep wells of mind stem from the twin pillars of the generative models of the brain and the embodied and situated nature of our cognitive systems. The depth of the human mind is not hidden away, nor a confabulation, but emerges in the rich interactive nature of the practices of mind in the world.

5. The real depth of the interactive mind

Chater argues that “Our brains are, then, relentless and compelling improvisers, creating the mind, moment by moment” (p. 220). We have argued this is only partly right. It is true we can be masterly improvisors, but these abilities emerge from our slowly acquisition of deep and world-revealing knowledge. The skillful mind is created and maintained over durations that are deeper and more extended than momentary improvisation. The case of learning to play cello has helped motivate and illustrate our proposal. It allowed us to present the means by which human beings acquire deep knowledge and the capacity for skilled, involved action. In the previous section, we have argued there are two sources of the depth of the human mind. The first source of depth is the hierarchical predictive knowledge we bring to those situations, as PP teaches us, our brains can be considered as a collection of highly integrated generative models or, more canonically, just one deeply-integrated generative model. The second source of depth is our rich and skillful embedding in the world. We are (often) skilled and profound manipulators of the environment we find ourselves in. The development of skillful performance can be understood as a gradual elaboration of skills, and with it a gradual expansion and refinement of our perceptual capabilities. By linking this account to predictive processing, we have also argued that these skills are grounded in the real causal structure of the world, and our abilities to intervene in it. As in our example of the child’s growing sensitivity to music and through their growing abilities to control musical performance, skills are slowly appropriated and mastery gradually developed. Learning to play the cello is a process of ever-deepening mastery but such skilled mastery applies across the range of our worldly interactions. Skills developed in one context can (often) be deployed in an increasing range of related scenarios, i.e., learning to play vibrato in a simple piece can later be applied to many other pieces of music. Thus, the novice is able to achieve more sophisticated and resilient skills of interpretation and performance.

However, one of the virtues of our proposal is that it also allows us to explain how and why mental depth can also be lost, or at least impeded, when our skillful practices become untied from their moorings, in the sorts of situations and contexts in which they were first developed. Situated skills are always relatively dependent upon context and as such are neither infinitely adaptive nor instantly redeployable, at least in short time horizons. Take, for instance, the reported experiences of teaching during the current COVID-19 pandemic via platforms such as Zoom or Google Meet. The experience of many educators living through the pandemic and attempting to rapidly convert teaching materials and practices to online platforms was experienced by many as a losing of capacities—being unable to easily interact with students, having that interaction denaturalized. We suggest that the embodied repertoire of skills that had been built after years of teaching in a classroom, were suddenly interrupted as many adept teachers and lecturers discovered they could not be easily put to use in the virtual classroom. Importantly, many such skills are typically in the background, relied upon implicitly. Habits such as walking across the class to capture the students' attention became inaccessible inhibiting the abilities of many experienced educators to engage with students in the same way as before. The new ecological setting of online learning thus can strip away skills that were previously taken for granted. When reflecting on our own experience and the many anecdotal reports, we believe that the dramatic change in ecological setting and resultant inhibition of relied upon skills should be described as an experience of loss of mental depth.²⁹ Such an experience is characterized by the feeling that abilities seem much more brittle or shallow and a characteristic fatigue in performing “the same” tasks that had previously seemed unproblematic via a new medium.

Prior to the confinement many educators lacked familiarization with practices for online teaching and learning, and that can also partly explain the feeling of losing depth and the diminished capacity for improvisation and fluent action. Of course, this is not to say it was impossible to teach well using tools like Zoom etc. In fact, as time passed by, we expect many educators would find themselves able to develop new refined practices responsive to online group dynamics and anchored in the new tools. That is why, after a sudden change in the usual situation, a period of training and adjustment was needed to regain the feeling of depth. The loss of the feeling of depth reveals, we think, something real. As presential teaching is replaced by online formats, many of our skillful practices become inhibited and new ones have to be developed. The reality of depth is revealed when it is lost and sometimes regained. We think this experience can be explained in virtue of the fact that much of the background processes upon which teaching relies: the situated and embodied context in which habits and skills are embedded were stripped away. Our skills always relate to, or at least are formed, in particular environments and our cognitive systems are intimately tuned to the material culture in which we are embedded.

²⁹ This might have been a collective experience. Lecturers in England have been reported as “feeling out of their depth.” Source: <https://www.theguardian.com/education/2020/dec/03/i-feel-out-of-my-depth-university-lecturers-in-england-on-the-impact-of-the-pandemic> (last accessed: January 29, 2021).

Summing up, our claim that mental depth is embodied and situated has allowed us to argue that depth is not illusory, but we should reconceive what both psychology and folk-psychology took mental depth to be.

Before concluding, we want to consider a possible objection to our account. One could argue, drawing from Chater's remarks, that despite our analysis of skillful practices the notion of situated depth still mental depth is apparent and confabulated. To allay this worry, let us finish by explaining why it's not the case.

There is really nothing obviously confabulated about skilled practices. We need to theoretically separate our skillful embodied interactions with tools and others, from stories we tell about them. It is important to see here that we are not denying that human minds do confabulate explanations, including the narrative explanation of themselves. But this "narrativity" is not the real source of our cognitive depth. Depth is much better located in skillful practices, including our "folk psychological" practices of interacting with others (Gallagher, 2001; McGeer, 2001; Zawidzki, 2008). Much social cognition may also be framed as a form of skilled practice and our cognitive depth tends to emerge in those selfsame skilled practices of interacting with others, often against the background of a rich environment. Such 'practices of mind' again are not illusory albeit our abilities to fully give verbal account of those practices (narrative) may be, under many conditions, rather weak. As the previous cases have suggested, human mental depth, both in its practiced reality and occasional absences is best understood not as a confabulation but as a finely constituted, deeply situated, and also sometimes precarious set of practices.

This supports our claim that the depth of the human mind is real but lies not in unconscious streams of thought that has been bequeathed to our folk-psychology by Freud but in our worldly practices. Chater (2018) is probably correct that there is no such thing as an unconscious stream of thought. Albeit, even here—as we have remarked—it is not quite true to say there is no background processing. The constant readjustment of the generative models of the brain in virtue of precision-estimation mechanisms, under predictive processing is indeed a form of background processing. It is the processing by which our sensitivities to the deep causal structure of the world are revealed to us through the ongoing history of our actions. This readjustment process, and indeed the deployment of active inference is a largely unconscious and background, or put another way, subpersonal, process. There is also a sense that the way it allows us to "see as", that is, respond to and represent the deep causal structure of the world, can be seen as a deep and hidden reservoir of much of our activities and mental lives.

6. Conclusion

We began this paper by introducing a recent challenge put forward by Chater (2018) to the idea of mental depth, an idea that deeply permeates our folk psychological image and plays an important role in contemporary psychology. Contrary to traditional psychology and our folk image, Chater (2018) argues that the mind is not deep and, consequently,

we are operating under an illusion about the nature of our own minds and cognitive processes. In this paper, we have put forward an account of mental depth that offers a partial reconciliation with the idea that we are deep mental creatures, but that at the same time is in tune with recent scientific approaches regarding our mental architecture. More precisely, we have partially agreed with Chater in that depth of the mind is not to be found (solely) on the deep background neuronal processing of the brain. Departing from Chater, we have claimed that this does not mean that our minds are not deep. To motivate our account, we have introduced a phenomenologically informed vignette that concerns the acquisition of mental depth through learning to play the cello. And we have relied on two different sources to explain this phenomenology. First, the predictive processing framework has allowed us to make a proposition about the brain's contribution to mental depth; specifically, in the instantiation of context-sensitive hierarchical knowledge produced by predictive processing systems. Second, we have drawn from a situated approach to cognition, to argue that the source of mental depth is located in our embodied skills and the situations in which we are embedded. Our moral is that mental depth is real, but we should move beyond its traditional understanding. Our mental depth is not the result of a deep unconscious or at least not only, but it results from our skilled engagement with a structured and sustained environment that we inhabit and act upon. Consequently, in order to understand human mental depth, we need not look only in the deep recesses of the brain but in how brains, bodies, and the world come together in skilled action. We find depth through worldly engagement and that is, for us, real enough.

Funding

Robert W. Clowes: This work has been supported by the Portuguese Foundation for Science and Technology (FCT) and CHIST-ERA Programme the via two projects: "EITNS – Emergence in the Natural Sciences: Towards a New Paradigm (PTDC/FER-HFC/30665/2017), and the Antidote project (CHIST-ERA/0002/2019).

Gloria Andrada: This work is funded by national funds through FCT—Fundação para a Ciência e a Tecnologia under the project UIDB/00183/2020.

Acknowledgements

We want to thank the organisers of the 2019 meeting in Porto (<http://avant.edu.pl/trends4>), especially Michał Piekarski and Witold Wachowski for their diligence. The paper grew from that the opportunity for detailed discussion and the interaction with those present at the meeting on "The Explanatory Power of Predictive Processing and Its Limits", especially Wanja Wiese, Nico Orlandi, Paul Thagard, Georg Theiner, Krzysztof Dolega, Steven Gouveia and Paweł Gładziejewski and we thank all those present at the workshop for their input. Special thanks to Marcin Miłkowski for earlier conversation and encouragement on these themes. Robert also presented some early versions of these ideas to the LMRG research in progress seminar and especially wants to thank Klaus

Gaertner, Nuno Venturinha and Paul Smart for their feedback. We also want to thank the two anonymous reviewers for detailed critiques of the paper and once again to Witold for his friendly guidance and forbearance through the publication process.

References

- Andrada, G. (2020). Transparency and the phenomenology of extended cognition. *LÍMITE Interdisciplinary Journal of Philosophy & Psychology*, 15(20), 1–17.
- Barrett, L. F., & Bar, M. (2009). See it with feeling: Affective predictions during object perception. *Philosophical Transactions of the Royal Society ,B*, 364(1521), 1325–1334.
- Blackmore, S. (2002). There is No Stream of Consciousness. *Journal of Consciousness Studies*, 9(5-6), 17-28.
- Bredlau, S. M. (2010) A Respectful World: Merleau-Ponty and the Experience of Depth, *Human Studies*, 33, 411–423.
- Chater, N. (2018). *The Mind Is Flat: The Remarkable Shallowness of the Improvising Brain*: Yale University Press.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: The MIT Press.
- Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Oxford: Oxford University Press.
- Clark, A. (2015). What ‘Extended Me’ knows. *Synthese*, 192(11), 3757–3775.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford: Oxford University Press.
- Clowes, R.W. (2014). Thinking in the Cloud: The Cognitive Incorporation of Cloud-Based Technology. *Philosophy and Technology*, 28(2): 261–296.
- Clowes, R.W. (2019). Immaterial Engagement: Human Agency and the Cognitive Ecology of the Internet. *Phenomenology and the Cognitive Sciences*, 18(1), 259–279.
- Clowes, R. W., & Gärtner, K. (2020). The Pre-Reflective Situational Self. *Topoi*, 623–637, <https://doi.org/10.1007/s11245-018-9598-5>.
- Dennett, D. C. (1991). *Consciousness explained*. Harmondsworth: Penguin Books.
- Dennett, D. C. (1992). The Self as a centre of Narrative Gravity. In F. Kessel, P. Cole & D. Johnson, (eds.), *Self and Consciousness: Multiple Perspectives*. Hillsdale, NJ: Erlbaum.
- Dennett, D. C. (2001). Surprise, surprise. *Behavioral and Brain Sciences*, 24(5), 982.
- Dołęga, K. & Dewhurst, J. E. (2021). Fame in the predictive brain: a deflationary approach to explaining consciousness in the prediction error minimization framework. *Synthese*, 198(8), 7781–7806.
- Donaldson, M. (1979). *Children's minds*. W. W. Norton & Company.

- Donaldson, M. (1992). *Human minds: An exploration*. Allen Lane/Viking Penguin.
- Fabry, R. (2018). Betwixt and between: the enculturated predictive processing approach to cognition. *Synthese*, 195(6), 2483–2518.
- Fodor, J. (1975). *The Language of Thought*. Hassocks, UK: Harvester Press.
- Frankish, K. (2016). Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11–12), 11–39.
- Gallagher, S. (2001). The Practice of Mind. *Journal of Consciousness Studies*, 8(5–7), 83–108.
- Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford: Oxford University Press.
- Gallagher, S. & Allen, M. (2018). Active inference, enactivism and the hermeneutics of social cognition. *Synthese* 195(6), 2627–2648.
- Gazzaniga, M. S. & LeDoux, J. (1978). *The Integrated Mind*. New York: Plenum.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.
- Hohwy, J. & Seth, A. (2020). Predictive processing as a systematic basis for identifying the neural correlates of consciousness. *Philosophy and the Mind Sciences*, 1(II).
- Hurley, S. (1998). *Consciousness in Action*. Cambridge MA: Harvard University Press.
- Hutchins, E. (1995) *Cognition in the Wild*. MIT Press.
- Hutchins, E. (2011). Enculturating the Supersized Mind. *Philosophical Studies*, 152(3), 437–446.
- Krueger, J. (2013). Affordances and the musically extended mind. *Frontiers in Psychology*, 4, 1–12.
- Mantel, H. (2011). *Wolf hall*. Fourth Estate (UK).
- McGeer, V. (2001). Psycho-practice, psycho-theory and the contrastive case of autism. How practices of mind become second-nature. *Journal of Consciousness Studies*, 5–7, 109–132.
- Menary, R. (2007). Cognitive integration: mind and cognition unbounded. Palgrave Macmillan.
- Menary, R. (2010). Dimensions of Mind, *Phenomenology and the Cognitive Sciences*, 9(4), 561–578.
- Menary, R. (2018). Cognitive Integration How Culture Transforms Us and Extends Our Cognitive Capabilities. In A. Newen, L. De Bruin, & S. Gallagher (eds.), *The Oxford handbook of 4E cognition* (pp. 187–215), Oxford: Oxford University Press.
- Newen, A., De Bruin, L., & Gallagher, S. (eds.). (2018). *The Oxford handbook of 4E cognition*: Oxford University Press.
- Noë, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9(5–6), 1–12.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: Bradford Books, MIT Press.

- O'Regan, J. K. (2002). The Grand Illusion and Sensorimotor Approaches to Consciousness. *Journal of Consciousness Studies*, 9(5–6).
- O'Regan, J. K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–973.
- Power, Michael J. & Brewin, C. R. (1991). From Freud to cognitive science: A contemporary account of the unconscious. *British Journal of Clinical Psychology*, 30, 289–310.
- Reed, E. (1993). The intention to use a specific affordance: a conceptual framework for psychology. In R. H. Wozniak & K. W. Fischer (eds.), *Development in Context: Activity and Thinking in Specific Environments*. Taylor and Francis Group.
- Rowlands, M. (1999). *The Body in Mind: Understanding Cognitive Processes*. Cambridge: Cambridge University Press.
- Rowlands, M. (2009). Extended Cognition and the Mark of the Cognitive. *Philosophical Psychology*, 22(1), 1–19.
- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1(7), 261–267.
- Schlicht, T. & Dolega, K. (2021). You can't always get what you want: Predictive processing and consciousness. *Philosophy and the Mind Sciences*, 2.
- Simons, D. J. & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9(1), 16–20.
- Sutton, J., Harris, C. B., Keil, P. G., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9 (4), 521–560.
- Suzuki, S. (1991). *Suzuki cello scholl, cello part 3* (Vol. 3). Alfred Music Publishing.
- Wester, D. (1996). The Scientific Status of Unconscious Processes: Is Freud Really Dead? *Journal of American Psychoanalytic Association*, 47(4), 1061–1106, <https://doi.org/10.1177/000306519904700404>.
- Zawidzki, T. W. (2008). The function of folk psychology: mind reading or mind shaping? *Philosophical Explorations*, 11(3), 193–210.