

Kristjan Loorits<sup>1</sup>

# *Qualities in the World, in Science, and in Consciousness*

**Abstract:** *It has been argued, most famously by David Chalmers, that all objects of the so-called traditional sciences (from physics to neuroscience) are analysable in structural terms, whereas consciousness has qualitative properties that are irreducibly non-structural. From that it has been concluded that consciousness cannot be explained by traditional sciences. Some illusionists have responded by proposing that the apparently non-structural features of consciousness are in fact fully structural and merely seem to be non-structural. I argue that such a position is tenable, but only if the non-structural ‘seemings’ are interpreted as perspectival phenomena and not as theorists’ fictions or absolute nothingness. The resulting view allows us to ignore sensory qualities in the context of natural sciences while acknowledging their importance in the moral and philosophical domains. The proposed perspectivist interpretation also provides an account of how sensory qualities can be observed and talked about despite having no autonomous causal powers.*

## **1. Qualities as Perspective-Induced Non-structural Impressions**

One of the deepest disagreements in the field of consciousness research concerns the explanatory sufficiency of the so-called

Correspondence:  
Email: [kristjan.loorits@helsinki.fi](mailto:kristjan.loorits@helsinki.fi)

---

<sup>1</sup> University of Helsinki, Finland.

traditional scientific framework and methods. That disagreement boils down to the question whether we need to or need not posit new fundamental laws, principles, or entities in order to explain consciousness.

The need for a new fundamental ingredient is often defended by appealing to the famous *hard problem of consciousness*. Arguably, scientific explanations are only suited to account for the *functions* or *structure and dynamics* of things, whereas consciousness has besides all its functional or structural-dynamical properties also *irreducibly qualitative* properties, which define *what it is like to be* conscious (e.g. Nagel, 1974; Chalmers, 1996; 2003). Accordingly, the traditional sciences (e.g. psychology and cognitive neuroscience) are in principle sufficient to explain all human behaviour (including the behavioural manifestations of consciousness), but not the distinctive *qualitative feel* of being conscious.

It has been therefore suggested that in order to solve the hard problem we must postulate the existence of a fundamental *phenomenal* or *proto-phenomenal* aspect of reality that *accompanies* some (or all) empirically accessible functional and structural phenomena in a law-like manner (e.g. Chalmers, 1996; Russell, 1927; Tononi and Koch, 2015; Strawson, 2006; Seager, 2012; Hameroff and Penrose, 2014; Atmanspacher, 2020). The nature of the hypothetical extra ingredient is conceived differently in different theories; being sometimes compared to paradigmatic fundamental phenomena like mass and charge (Tononi and Koch, 2015; Hameroff and Penrose, 2014) or seen as the underlying basis (or the internal nature) of some (or all) empirically accessible features of reality (e.g. Russell, 1927; Chalmers, 1996; 2003; Strawson, 2006; Seager, 2012). These approaches are sometimes lumped together under the label of *panpsychism* (especially by their critics, e.g. Seth, 2021), but since not all of them insist that consciousness is ubiquitous, it might be more appropriate to call them *fundamental phenomenality* approaches.

One of the biggest problems for the fundamental phenomenality view is a threat of *epiphenomenalism*: if it is indeed the case that the traditional sciences (i.e. the sciences that do not recognize fundamental phenomenality or proto-phenomenality) are in principle capable of explaining all human behaviour — including the behavioural manifestations of consciousness — then the fundamental phenomenality (or proto-phenomenality) would have no influence on our behaviour, including our physically expressed arguments for the existence of fundamental phenomenality. Thus, the view implies that all physically expressed (written, spoken, etc.) arguments that support

the fundamental phenomenality view would be put forward even if the view were not true. Although the epiphenomenalist implications do not render the fundamental phenomenality view logically inconsistent or impossible to hold (for a recent defence of epiphenomenalism, see Robinson, 2019), they make it difficult to defend the view in a credible manner. Every attempt to do so would contain an implicit claim that such an attempt would be made even if the view were incorrect. It is significant that there are no similar *expression problems* with any other fundamental entities recognized by science.

The above sketch is very rough but helps to outline the general nature of the situation: while the hard problem pushes us towards the fundamental phenomenality view, the threat of epiphenomenalism and the resulting expression problem pulls us away from it. And although there are numerous arguments both for and against taking either direction, none have succeeded settling the issue or bringing forth a substantial consensus.

The position taken and defended in this paper is that there is no need to posit any ‘new’ fundamental entities or aspects of reality in order to explain consciousness. Instead, it is proposed that the apparent need for such a special ingredient may be satisfied by acknowledging an ‘old’ element that can be found in many domains of life, including scientific practices — an element whose existence is implicitly recognized by every scientist yet systematically overlooked in philosophy of science.

The overlooked ingredient may be called by different names depending on the context and domain, but for an initial intuitive grasp, let us call it simply quality. The notion of quality in question conforms to how the term is used in the context of the hard problem where the qualitative properties are contrasted with the structural or functional ones (e.g. Chalmers, 1995; 1996). Following that use of the term, qualities can be understood as properties that are at least *prima facie* non-structural.

Aside from consciousness, there are many other domains where we find properties we can easily name and recognize without being able to provide structural analyses of their nature. In a daily context, we can talk about the cuteness of kittens or the charisma of politicians without knowing how to analyse those properties in fully structural terms. Yet, we are normally not puzzled by the non-structural aspect of cuteness or charisma. Instead, we are ready to accept that although these properties may *seem* not fully structural from our everyday perspective, they are probably structurally analysable in principle (and

are therefore fully structural by their ‘true’ perspective-independent nature). In other words, we are quite happy with the idea that the structurally unanalysable aspect of these properties is nothing but a perspective-induced impression. Accordingly, qualitative differences between cute and ugly or charismatic and awkward would be ‘genuinely’ qualitative only from a perspective that provides no access to the fine-grained relational structures that we assume to underlie these differences.

In most cases, these non-structural qualitative impressions do not provoke any deeper philosophical questions because it is understood how changing perspective would allow us, at least in principle, to eliminate their non-structural appearance. Just as these non-structural qualities ‘arise’ when we adopt perspectives that prevent us from discerning some structural details, so do they ‘dissolve’ when we adopt perspectives that allow us to see those structural details with sufficient clarity.

In the domain of science, the above notion of quality corresponds to the distinctive vagueness of poorly understood preliminary ideas. Vague and imprecise ideas are not only entertained privately in the minds of individual scientists but are often communicated and discussed (admittedly, with a high risk of miscommunication). More importantly, such ideas are *indispensable* during the phases where the corresponding precise formulations are yet to be developed. Science does not progress by jumping mechanically from one precise structural description to the next, from one experiment to another, from one well-formulated theory to another. Besides carefully conducted experiments and precisely formulated theories there are also gut feelings, vague hypotheses, and imprecise preliminary ideas. Most importantly, those vague preliminary ideas are not just structurally simpler or less accurate versions of the precise *structural* accounts they may ultimately lead to. Taken *as perspectival* phenomena (i.e. as they appear *from the perspectives of those who have them*), the genuinely vague ideas are simply not fully structural. Although most vague ideas have some recognizable structure (just as our everyday impressions of cuteness or charisma), their very vagueness consists in the fact that the identifiable structural components do not exhaust their nature. The additional non-structural element can be *observed* and *attended* to — as we do whenever we acknowledge that we have a vague intuition or a gut feeling whose content is partly ‘ineffable’ in the sense that we fail to articulate it in precise structural terms. Moreover, the vague and ‘ineffable’ part is quite often the one we focus our

attention on when trying to ‘make the idea clearer’ — as if knowing that the missing structural details are somehow hidden there and can be brought to light through our patient attentive efforts.

However, when vague ideas are attended, shared, and discussed, scientists are usually not interested in their peculiar non-structural character (at least not in the way philosophers are interested in the non-structural character of qualia). Instead, scientists are interested in finding and adopting perspectives that would allow them to get rid of that vagueness. And once they succeed, the non-structural element appears to be eliminated — leaving nothing to be philosophically interested about.

One of the main proposals of this paper is that the non-structural aspect of our vague ideas and the infamous sensory qualia are perspective-induced phenomena of the same basic kind. Both are marked by their ostensibly non-structural qualitative character that can only be observed from the perspective that induces it. In the domain of science, the psychologically satisfying sense of *understanding* is often achieved by learning to switch perspectives — moving in one’s mind back and forth between a vague and qualitative grasp and the corresponding precise structural description. When it comes to consciousness, we are normally incapable of abandoning the perspective from which qualia appear non-structural. Yet, even here it is possible to evoke minor shifts in perspective with a result of revealing some of the structures that form the perspective-independent basis of qualia (more on that in the following section). And conversely, even if we assume that all objects of natural sciences can be analysed in structural terms in principle, for each of us there are plenty of objects we can easily recognize yet fail to analyse in fully structural terms. Based on that it can be argued that the difference between qualia and the non-structural aspect of our vague ideas is more in degree than in kind: qualia are simply more stable and harder to ‘structuralize’.

Even if we gained detailed scientific knowledge about the underlying structures of qualia, our cognitive architecture would still preserve the perspective that displays most of those structures in a qualitative manner. Thus, in the case of qualia, the psychologically satisfying sense of understanding can be never fully attained — the perspective-induced sensory qualities cannot be eliminated because the perspectives that uphold them cannot be entirely abandoned. The reasons for that are probably multiple and complex, but one of them might be that qualia provide an organism with a vast amount of tacit yet instantly available information about the environment (as proposed

by various very different theories, e.g. Koch, 2004; Beaton, 2013; Dennett, 1988; 1991; 2015) and replacing that fast and rich mode of information with a more structured and explicit but slower one would put an organism at risk (e.g. Koch, 2004, pp. 231–47).

However, this is only half the story. The above interpretation raises the question as to how and if the broadly structuralist scientific worldview can incorporate the non-structural perspective-dependent qualities (and as argued in the next section, it neither can nor needs to). Setting sensory qualia aside, we may say that scientists are experts at ‘eliminating’ the perspective-induced non-structural phenomena by abandoning the perspectives that uphold them. But that kind of ‘elimination’ does not amount to establishing that there are no *perspectival* non-structural phenomena. Whenever we acknowledge that something *seems* (or seemed) to be qualitative or that we have (or had) an idea (intuition of a gut feeling) without being able to grasp and articulate it in a precise structural manner, we implicitly admit that some perspectival phenomena (the way things seem(ed) or appear(ed) to us) are not fully structural. And in a way, that understanding conforms to some intuitions behind the fundamental phenomenality view: if science is in the business of discovering patterns in the data (e.g. Ladyman *et al.*, 2007) and revealing the structure and dynamics of the so-called physical objects and processes (Chalmers, 2003), then the existence of non-structural perspectival phenomena is, from the scientific point of view, a surprising anomaly — something not fully captured by the structural description of the world.

It is important to acknowledge that when it comes to perspective-induced non-structural phenomena (either qualia proper or qualitative elements in our vague ideas, unanalysed impressions, and the like), we are not dealing with merely unobserved or unnoticed structures but with a distinctively non-structural content that can be observed and attended to. When I look at the bright summer sky and am baffled by my failure to analyse the experienced blueness in structural terms, I am aware of *something* that is not fully structural from my perspective. The same is true about the vague and inarticulate (or ‘ineffable’) part of my gut feeling I can focus my attention on without succeeding to formulate it in structural terms. At the same time there are myriad structures I am completely unaware of and whose existence is entirely irrelevant when it comes to describing how things seem to me. Those are the structures that are from my perspective *nothing*, while sensory qualities and vague impressions are from my perspective *something* I fail to analyse in fully structural terms. Thus,

even if all my non-structural impressions correspond to some misperceived or poorly discerned structures (in the sense specified in the following section), their non-structural character is nevertheless instantiated as a property of *how they seem to me*.

Despite its realist undertones, the above interpretation comes rather close to a version of *illusionism* about qualia (without sharing its eliminativist inclinations). It has been argued, among many others, by Dennett (1988; 1991), Markkula (2015), and via a slightly different route (and without an explicit commitment to the label of illusionism) by Crick and Koch (1998; 2003), Koch (2004), O'Regan (2011), Beaton (2013), and Loorits (2014; 2019) that qualia (or what is usually meant by the term) are in fact complex structural phenomena that merely *seem* to be non-structural and irreducibly qualitative (the position examined more closely in the following section). Some illusionists, most famously Dennett (e.g. 1988; 1991), are even against using the very notion of qualia, but in this paper the notion is used, for the sake of simplicity, also in the illusionist context (in the sense of what *seems* to be non-structural in our experiences).

The general reasoning behind the above form of illusionism is that if qualia are *really* structural and merely *seem* to be non-structural, then their puzzling non-structural character no longer requires an explanation — after all, then qualia are not *really* non-structural. By analogy, although a physics student's initial grasp of, say, special relativity may be partially vague, there is no need to explain the non-structural vagueness of special relativity — the content of special relativity is not *really* vague (at least not in the sense it appears to the student). More generally, once we acknowledge that the perspective-induced vagueness occurs in many other domains without undermining the scientific structuralist worldview, the perspective-induced qualitative-ness of qualia can be reconciled with the scientific structuralism in a similar manner. But it does not follow that the existence of perspective-induced non-structural qualities can be flat out denied or declared philosophically uninteresting.

What makes perspective-induced qualities interesting is that they are easy to observe and talk about (as we routinely do when something seems vague and non-structural to us), yet hard to analyse in scientific terms. Science strives towards maximally precise structural descriptions (even when those descriptions concern only hypotheses or approximations), and thus the phenomena that are essentially vague and non-structural escape its descriptive scope. In light of the illusionist reasoning, science is free to *dismiss* those non-structural phenom-

ena as perspective-induced impressions or mere seemings. However, such a dismissal does not amount to genuine explanation (nor to an absolute elimination). Rather, it expresses an attitude according to which perspective-induced non-structural artefacts (including sensory qualia) are something that science is not *supposed to* explain (more on that in the following section).

Following the above interpretation, the position defended in this paper rejects the idea that qualia belong to the perspective-independent domain (among mass, charge, and other fundamental phenomena recognized by science). On the other hand, the position is also in tension with the conventional understanding of physicalism (or what is usually assumed under that label). Therefore, let us see next how the perspectivist interpretation of qualia can be clarified against the background of these two more influential alternatives.

## 2. Qualia and Physicalism

As the idea of *reductive physicalism* (according to which the objects of special sciences can be ‘reduced away’ to the underlying micro-physical processes) has become increasingly unpopular, let us try to capture the spirit of physicalism in less controversial terms of *reductive explanations*. Put roughly, a phenomenon is reductively explainable or simply *fully* explainable in terms of other phenomena if and only if an appropriate story about the latter contains (at least implicitly) a description of the former (and in the case of multiply realizable objects, the pool of explanatorily feasible ‘lower-level’ phenomena may be indefinitely large and diverse).

When it comes to physicalism about consciousness, most contemporary proponents of the view are not particularly interested in defending physicalism as a general metaphysical thesis. Instead, the term is applied rather loosely to different approaches sharing the basic idea that consciousness, with all its puzzling features (such as phenomenality and qualia), can be *in principle* reductively explained (or fully explained or fully analysed) in terms of some ‘ordinary’ phenomena that are not puzzling in that sense (e.g. brain processes, behavioural capacities, dispositions, information processing, representations, affordances, etc.).

In light of the above loose notion of physicalism, the challenge posed by the hard problem (in line with its original formulation by Chalmers, 1995; 1996) is that all the ‘ordinary’ phenomena that are viable candidates for the explanatory basis of consciousness can be



analysed in structural terms while the qualitative character of consciousness appears to be non-structural. For how could any story about structure and relations *contain* (even implicitly) a description of a non-structural quality? What could we even mean by a *description* of a non-structural quality? Are not descriptions, at least in the context of science, specifications of structure and relations?

As noted in the previous section, a way to avoid the above kind of questions (and to address the hard problem) is to deny the non-structural nature of qualia. One of the most interesting and best-known attempts to reveal the structural nature of our apparently simple qualitative impressions is made by Daniel Dennett (e.g. 1988; 1991). A particularly powerful example from Dennett's repertoire concerns the auditory experience of a low guitar sound (1988; 1991, pp. 49–50). Presumably, the sound appears to an untrained ear just as primitive and unanalysable as a typical colour quale (e.g. experienced blueness) — defined solely by its ineffable 'guitarish' character. And yet, with the help of simple exercises, the subject can learn to distinguish (some of) the individual overtones of the sound and realize that her seemingly unanalysable and primitive auditory quale has a subtle internal structure after all.

It is easy to find similar examples about other sensory modalities. As Dennett (2015, p. 8) emphasizes (referring to the work of Huron, 2006), many specialists, such as wine tasters and musicians, can learn to distinguish the 'combining elements' of certain qualitative experiences that 'used to seem atomic and unanalyzable'. For instance, an apparently primitive gustatory quale (an experienced flavour) may be suddenly perceived as a bundle of flavours of individual ingredients. The significance of such examples is that they help to demonstrate how and in what sense a subject can be mistaken — and *admit* of having been mistaken — in thinking that a particular sensory quality is non-structural and primitive.

It is reasonable to assume that the subtle structures revealed by such disillusioning exercises correspond to some empirically accessible neurocognitive structures. For instance, according to Crick and Koch's neurobiological theory (Crick and Koch, 1998; 2003; Koch, 2004), qualia are vast networks of non-conscious associations. Thus, the distinctive character of an experienced redness would be 'composed' of a vast number of non-conscious associations with different red objects (e.g. ripe tomatoes, strawberries, blood, sunsets, the subject's first car, etc.) and the associated events and behavioural strategies (e.g. eating strawberries and avoiding blood) (Koch, 2004, pp. 233–5).

At the neurobiological level, the experience of redness would be ‘encoded’ by an increased activity in a small group of neurons, the so-called *essential node* that ‘represents’ redness to the rest of the brain. And each of the unconscious associations (that contribute to the qualitative character of redness) would correspond to a direct mono-synaptic axonal connection between the essential node for redness and the essential node for some associated content (e.g. strawberries or sunsets). Every essential node would also project directly (via mono-synaptic axonal connection) to the planning modules of the brain — that is why the content of consciousness is normally reportable. So, the claim that the qualia-forming associations are non-conscious would correspond to the neurobiological claim that the activity in the individual essential nodes that encode the associated content is too weak to have a clearly discernable effect on the planning modules (that is why we cannot report the content of those individual non-conscious associations — neither to ourselves nor to others).

The above sketch is obviously very rough (for the most comprehensive account see Koch, 2004) but sufficient for the present purposes. One of the aims of Dennett’s guitar sound example was to make it intuitively easier to accept that the apparently atomic sensory qualities are in fact structural. Crick and Koch’s theory allows us to support that project with a further thought experiment. As I have argued elsewhere (Loorits, 2014; 2019), if qualia are structural in a manner Crick and Koch propose, then it should be possible, at least in principle, to experience the structural nature of any quale we choose. That would require, first, locating the essential nodes that are responsible for some key non-conscious components of the quale in question, and second, increasing the neural activity in those essential nodes. The expected results should be similar to the structure-revealing experience of a guitar sound. Although the subject would not re-experience the entire quale in a fully structural manner, she would experience some of its previously hidden components and realize that she was wrong in believing that the quale is primitive and unanalysable. And although the revealed components of the original quale have qualitative characters of their own, the subject would understand that the structure-revealing procedure could be also applied to each of these. In short, the subject would now have a phenomenological basis to accept the possibility that none of her qualia are irreducibly non-structural.

In light of the above considerations, it appears that the ostensibly crucial difference between sensory qualia and other perspective-

dependent qualities is not as clear-cut as usually assumed. Although other qualities may be easier to structuralize, it is possible that sensory qualia can be structuralized in the same basic way. And even if no subject could ever experience her entire phenomenal world in a fully structural manner, neither could any scientist grasp her entire domain of expertise in a clear and fully structural manner all at once. Figuratively speaking, in both cases one lives in a world that is largely qualitative and learns about its underlying structural nature by focusing her attentive and analytical gaze on one small bit at a time. And once one learns that the thereby revealed structures fit together and form a coherent whole, one is led towards a structuralist understanding of reality. And yet, none of the above suggests that the qualitative part of one's phenomenal world could be ever fully eliminated.

When aiming at maximally precise descriptions (e.g. mathematically formulated theories) that lay ground for scientific progress, it has proven to be a remarkably good strategy to dismiss perspective-dependent non-structural qualities as mere seemings and to abandon them (if possible) by changing perspectives. But besides this well-justified yet purely practical dismissal, there is no deeper metaphysical sense in which these non-structural qualities could be considered unreal (one might be tempted to argue that perspectival qualities are also causally impotent, but in the present context that would only mean that the precise structural accounts of relevant causal factors make the more vague and 'qualitative' accounts of those same factors redundant and dismissible). While the demand for objectivity in the sense of perspective-independence has led philosophers of science towards scientific structuralism, it would be misleading to say that science has taught us that *everything* is structural. Rather, *philosophy* of science has taught us that the distinction between 'real' and 'unreal' that is based on the notion of objectivity as perspective-independence leads to structuralism. It can be argued that even the most radical forms of scientific structuralism (e.g. the ontic structural realism promoted by Ladyman *et al.*, 2007) are not meant to apply to all perspectival phenomena (the way things seem). On the contrary, part of the reason why those radical forms of scientific structuralism are considered interesting positions is precisely that they are counter-intuitive and go against the way things seem.

To conclude, as far as the general notion of physicalism is committed to the idea that the qualitative character of consciousness is reductively explainable or fully eliminable, the perspectivist view defended in this paper is not a physicalist position. When it comes to

perspective-induced non-structural qualities, science has never aimed at reductive explanations but a sort of practical dismissal that is well justified based on the history of science. And even the most impressive cases of ‘structuralizing’ qualia do not provide a counter-example to that.

When a subject learns to discern the previously unregistered overtones, her experience changes dramatically. In the words of Dennett himself (1991, pp. 49–50): ‘The difference in experience is striking, but the complexity newly apprehended... was there all along.’ Since the complexity was there all along, what could possibly constitute the *striking difference* in experience? Quite simply, the difference is in how things *seem* before and after. And as emphasized in the previous section, the way things seemed before (when the complexity was not ‘apprehended’) cannot be characterized as a mere failure to attend to certain structures (although that failure is part of the story). From the perspective of the subject, the undiscerned overtones are not simply absent but have merged into an ‘ineffable’ non-structural impression she can easily observe and attend to. Thus, the disillusioning experience of individual overtones may help one to accept a structural account of her ‘guitarish’ quale, but such an account would never contain a description of the non-structural character that is experienced by a subject who *fails* to distinguish the individual overtones of the sound — simply because no structural account could ever contain a description of anything non-structural.

But even if we accept that the perspective-induced qualities are reductively inexplicable as outlined above, it does not follow that they are fundamental in the sense of requiring any special attention or recognition in the context of fundamental physics (or any natural sciences, for that matter).

### 3. Qualia and the Fundamental Phenomenality View

Most theories that have adopted the idea of fundamental phenomenality (or proto-phenomenality) associate phenomenality with some recognized or hypothesized empirical phenomena. For instance, according to the *integrated information theory* (e.g. Tononi and Koch, 2015; Tononi *et al.*, 2016; Oizumi, Albantakis and Tononi, 2014), ‘consciousness is a fundamental property possessed by physical systems having specific causal properties’ (Tononi and Koch, 2015, p. 1). According to the *orchestrated objective reduction theory*

developed by Hameroff and Penrose (e.g. 2014, pp. 70–1), ‘each action of [objective state reduction of superposed quantum states] is accompanied by a moment of *proto-consciousness*...’ Similarly, Pykkänen (2018) suggests, based on the Bohmian quantum-mechanical notion of *active information* (e.g. Bohm, 1989; Bohm and Hiley, 1993), that the *content* of active information might be proto-phenomenal or proto-qualitative.

While the above accounts link qualia and phenomenality to rather specific empirical phenomena, some others propose more general and basic connections between empirically accessible and phenomenally experienced domains. For instance, according to Chalmers (1996; 2003), phenomenality (or proto-phenomenality) is a fundamental feature of the world that *underlies* the physical reality (in his 1996 book, Chalmers proposed that the *structural* features of our phenomenal experiences correspond to the informational structures in the physical domain, but that the primitive phenomenal qualities are linked to the physical world by fundamental psychophysical laws). Even more broadly, while defending his position known as *reflexive monism*, Velmans (2009, p. 324, endnote 6) admits of assuming that ‘it is simply a “natural” empirical fact about the world that certain physical events in the brain (the neural correlates of consciousness) are accompanied by experiences’. By the same token, Velmans suggests that the relationship between phenomenal and neural states ‘follows some natural law, however mysterious this presently seems’ (*ibid.*).

Although Chalmers’ and Velmans’ approaches have many interesting similarities with the perspectivist interpretations defended in this paper, both depart from the latter by emphasizing the need for some fundamental psychophysical laws that would specify how the phenomenal and physical domains are related. In contrast, the perspectivist interpretation sees the relationship between phenomenal and physical as a particular form of the more general and widely acknowledged relationship between perspective-independent and perspective-dependent phenomena.

At this point it may be worth clarifying that for present purposes it is not necessary to assume that the distinction between perspective-dependence and perspective-independence is always clear-cut and has a uniform expression in all domains from daily life to quantum physics. Even if we assume that the distinction is always to some extent contextual or relative (e.g. Hautamäki, 2020), it is nevertheless meaningful in the contexts where we distinguish between reality and

appearance and where we disregard something as a mere seeming. And in those contexts, the relationship between perspectival phenomena and their perspective-independent bases is not the kind we would normally call fundamental (as demonstrated shortly below).

As emphasized in the first section, the perspectivist interpretation of qualia allows us to approach the most puzzling features of consciousness by comparing them with less puzzling perspectival phenomena. Accordingly, at one end of the spectrum we would find fully reducible perspective-induced *structural* phenomena, such as two-dimensional projections of three-dimensional objects or optical illusions that display some geometrical properties of physical objects in a distorted manner (e.g. a straight stick that appears to be bent in water). A step further we find somewhat vague and imprecise pre-scientific ideas that can be clarified (or structuralized) by changing perspectives. While the distinctive non-structural vagueness of these ideas cannot be reductively explained, it can be ‘eliminated’ by abandoning the perspectives that maintain it.

Further along the spectrum we find artistic expressions that are partly appreciated precisely because their meanings or content cannot be fully analysed in precise structural terms. The main difference between the vagueness in our pre-scientific ideas and in art lies in our attitude. While the vagueness in science is seen as an obstacle to be removed, in art the appropriate balance between clarity and vagueness is usually highly valued. But in neither case are we tempted to think that the element of vagueness is a fundamental feature of reality that must be connected to the empirically accessible structures by some unique and basic psychophysical or ‘qualistructural’ laws. Instead, we are happy to accept that the vagueness (either valuable or unwelcomed) is just a perspectival artefact that deserves no attention in the context of fundamental physics.

What sets qualia apart is their remarkably stable and ‘intense’ nature. While in the domains of art and science the non-structural vagueness and structural clarity are in constant flux — some ideas becoming clearer and more structuralized while others fade into unarticulated vague grasps — the paradigmatic sensory qualia are exceptionally difficult to structuralize (although not entirely impossible, as we saw in the previous section). Qualia do not ‘break’ easily under our attentive gaze, which is why we do not normally think of them as vague (we expect vague things to have a potential to become clearer, and in the case of qualia, that potential is usually not evident to us). However, philosophically problematic non-structural qualities

do not arise at the end of the spectrum where we find sensory qualia. They arise at the point where we, as structure-detecting creatures, lose track of some structural details due to the poor resolution of our ‘discernment machinery’ (to use an expression from Dennett, 2015). And even if we assume that such machinery itself and the structure it fails to discern (as well as the causal interactions between these two) can be given a precise structural description, the fact remains that such a description would not contain or capture the distinctive non-structural character of the perspective-induced vagueness (as no structural description ever could).

By focusing on these simple instances — the cases in which some structural details become ‘blurry’ from a particular viewpoint — it is easy to accept that emerging non-structural qualities are perspective-induced artefacts with no autonomous causal powers. And that brings us to another crucial difference between the fundamental phenomenality approaches and the perspectivist interpretation defended in this paper.

As mentioned at the beginning of this paper, one of the most general and difficult problems for the fundamental phenomenality view is a threat of epiphenomenalism and the resulting expression problem. If the phenomenal character of consciousness is a fundamental aspect of reality that is immediately experienced yet not needed to explain our behaviour, then it becomes a mystery how anybody could *express* the existence of such an immediately experienced phenomenality. The perspectivist interpretation of qualia avoids that problem by acknowledging an *innocent* form of ‘epiphenomenalism’ that is widely accepted and deemed uncontroversial: the causal redundancy of perspectival artefacts.

Clearly, we can observe and talk about all kinds of perspectival phenomena without referring to or even knowing their perspective-independent bases (e.g. rainbows, stage magic, optical illusions, etc.). And according to the conventional understanding, in all such cases the explanations of our expressive behaviour can be given, at least in principle, in terms of some perspective-independent processes. Put simply, perspectival artefacts are taken to have no autonomous causal powers even though they can be observed and talked about. To be sure, we may include perspectival phenomena (as virtually anything, see Menzies, 2007) in our causal stories, but what justifies the causal redundancy claim is the fact that we do not *have to* include them — the relevant perspective-independent processes are causally sufficient by themselves.

It is important to notice that there is nothing controversial or speculative about the above kind of causal redundancy. It is implicitly assumed in all the commonplace situations in which we have psychological or behavioural responses to some perspective-induced content without assuming that such a content carries any autonomous and irreducible causal powers. None of the versions of the fundamental phenomenality view provides us with as simple and illuminating an account of a *non-causal* relationship between physical (i.e. empirically accessible) and phenomenal that is nevertheless compatible with the existence of *expressible* observations and reports of phenomenal qualities.

While all versions of the fundamental phenomenality view are more or less speculative and ask us to bite rather large bullets (e.g. unique kinds of psychophysical laws that are experimentally undetectable), the perspectivist interpretation of qualia is based on observations and relies entirely on a distinction that is implicitly acknowledged by all scientists and philosophers. Regardless of one's position on scientific structuralism, everyone agrees that there are things that *seem* to be non-structural. If that were not the case, then all of us would be thoroughly structuralist from birth and even the purely conceptual distinction between structural and non-structural would make no sense.

The appeal of the fundamental phenomenality view may lie partially in its ostensibly simple and clear metaphysical message: phenomenal properties are irreducible and reductively unexplainable, and the best we can do is to track their robust and fundamental connections to some empirically accessible objective phenomena. In comparison, the perspectivist interpretation may seem tentative and evasive. Yet, its core idea is no less metaphysically clear. Accordingly, phenomenal properties are indeed irreducible and reductively unexplainable, but their relationship with the empirically accessible phenomena is not unique but the kind we already know from many other domains.

The fundamental phenomenality view may also seem to take consciousness more seriously by declaring that phenomenal properties are real and not illusory. But when it comes to such foundational metaphysical claims, it is necessary to specify what is meant by *being real* in a given context. Perspectivists may say that qualia are real in the sense of being observable and reportable and illusory in the sense of being perspectival artefacts with no counterpart in the perspective-independent domain and no autonomous causal powers — or something along those lines. And if the meaning of 'being real' needs to be



clarified in some further terms anyway, then it might be easier to bypass the reality-talk altogether and convey the intended meta-physical message directly in those clarifying terms: qualia are observable and reportable yet causally redundant perspectival artefacts.

Admittedly, many versions of the fundamental phenomenality view contain empirical hypotheses that may turn out to be true even if the fundamental phenomenality thesis itself is not. That is particularly clear in the cases of quantum consciousness theories (e.g. Hameroff and Penrose, 2014; Pylykkänen, 2018) and the integrated information theory (e.g. Tononi and Koch, 2015; Tononi *et al.*, 2016; Oizumi, Albantakis and Tononi, 2014). The plausibility of the empirical claims of these theories falls outside the scope of this paper, but as far as those theories are motivated by finding a place for the fundamental phenomenality in the perspective-independent domain, the perspectivist interpretation of qualia casts serious doubts on their prospects.

#### 4. Final Remarks

Although the label ‘hard problem’ is relatively new (appearing first in print in Chalmers, 1995), the challenges of explaining consciousness in scientific terms have preoccupied philosophers at least since the times of Descartes and Leibniz. Contemporary understanding of those challenges is influenced greatly by the classic paper by Thomas Nagel (1974), who argues famously that no amount of knowledge about the neurobiological structure and functions of bats’ brains would provide us with the knowledge of what it is like to be a bat. While highlighting the stark contrast between perspective-dependent (i.e. only *subjectively* accessible) phenomenal properties and the scientifically accessible perspective-independent (or objective) brain processes, Nagel failed to acknowledge that those contrasting categories are at the extreme ends of a more nuanced spectrum with a philosophically interesting middle part (as outlined in the previous section).

In light of the perspectivist interpretation defended in this paper, Chalmers’ formulation of the hard problem constitutes a significant and necessary step forward. His realization (inspired by Russell, 1927) that the problem originates from the tension between scientific structuralism and the non-structural character of sensory qualia allows us to address the problem as proposed in this paper. As noted in the previous section, we may now recognize that the transition from

structural to non-structural first occurs where some structural details merge together from certain perspectives. That way puzzling sensory qualia become comparable to other non-structural phenomena whose existence is (at least implicitly) acknowledged by all scientific structuralists.

Another line of development that prepares the ground for the position here defended is the emergence of empirically plausible structural accounts of qualia (e.g. Dennett, 1988; 1991; Koch, 2004; O'Regan, 2011; Beaton, 2013). Particularly significant are the examples of 'top-down' phenomenological structuralizations of qualia (e.g. Dennett, 1988; 1991) that can be backed up by bottom-up neurobiological hypotheses (e.g. Koch 2004). Without these phenomenologically illuminating examples, the claim that qualia are structural and merely seem to be non-structural would be difficult to understand and even more so to accept. Through these demonstrations their structural nature becomes intelligible; and consequently, their non-structural *appearance* becomes comparable with the perspective-dependent qualities in many other domains.

In sum, although the perspectivist interpretation of the non-structural character of qualia may strike one as a textbook 'armchair' proposal, it is in fact a position that can only be properly defended in light of quite recent empirical (or empirically informed) theories. In a way, we may think of the perspectivist interpretation as a philosophical 'plug-in' to empirical theories that try to explain sensory qualities in structural terms. Among these is the repeatedly mentioned theory of Crick and Koch (1998; 2003; Koch, 2004), which provides not only the neurobiological interpretation of the underlying structures of qualia, but also a suggestion of what could constitute a *viewpoint* from where qualia appear to be non-structural. As the theory proposes, only the content with a direct access (via monosynaptic axonal connections) to the planning modules (mainly in the prefrontal and anterior cingulate cortices) can become conscious. Thus, as Koch (2004) puts it, the planning modules of the brain constitute a sort of *non-conscious homunculus* — a structure that is not conscious in itself or of itself but only of the information it receives from the back of the cortex. Such an 'homunculus' would thereby correspond to the *viewpoint* that displays certain structural details in a qualitative manner.

Another approach that resonates well with the perspectivist interpretation of qualities is the attempt to understand consciousness in terms of Bayesian expectations (e.g. Clark, 2013; Clark, Friston and Wilkinson, 2019; Dennett, 2015; Friston, Wiese and Hobson, 2020).

In that context, Bayesian expectations concern the probable outcomes of different available actions. From the objective perspective-independent ‘viewpoint’, the functioning of the systems that are responsible for (planning) the actions and the content of their expectations can be given a precise structural description (in well-defined probabilistic terms), but from the perspective of those systems themselves, the content of their expectations is somewhat vague and not fully structural. How do we know that? Well, according to the view, we are such systems — and obviously we do not experience our sensory qualia as complex yet well-defined structures of probabilistic expectations. As noted earlier, according to Dennett (2015), the reason for that lies in the poor resolution of our discernment machinery.

As argued in the previous section, that is precisely what the perspectivist interpretation implies: qualities emerge and dissolve when some structural details are either lost or discerned from a certain perspective. What is nevertheless interesting and surprising is that the lost structural details do not, from the perspective of the perceiver, fade into nothingness but merge into vague qualitative impressions that can be observed and attended to. Dennett insists that there is no need to *project* these qualitative impressions into any kind of ‘subjective space’. I suppose he is right — just as there is no need to project any other perspectival phenomena into special metaphysical spaces. But what seems to confuse many of Dennett’s readers is his unnecessarily strong and provocative eliminativist rhetoric. Once we acknowledge that the puzzling qualitative character of consciousness is a perspectival artefact — a mere seeming — why should we insist further that it is nothing but ‘a theorist’s fiction’ (Dennett, 1988; 1991)? To call *observable* perspectival phenomena fictitious does not seem to capture their nature accurately.

Obviously, there would be a lot more to say about the possible compatibility and mutual relevance of the perspectivist interpretation of qualia and different consciousness theories. The aim of the above sketch was mainly to illustrate that the perspectivist interpretation would enhance our understanding of consciousness only if considered in the context of a theory that offers some structural account of the perspective-independent basis of qualia. Presumably, if one rejects all these theories and the general approach they represent, then one has no good reason to accept the perspectivist interpretation of qualia.

Although the focus of this paper has been almost exclusively on the hard problem and the qualitative character of consciousness, the perspectivist interpretation of non-structural qualities may have

interesting implications in many other areas. For instance, as already mentioned in passing, the content of art may be considered reductively unexplainable for the same reason as the non-structural character of qualia. But it could be argued further that sometimes the exploratory scope of art may also exceed the scope of science. In the case of particularly complex (e.g. psychological, social, political) phenomena, our understanding tends to remain always somewhat vague and intuitive while even the fragmentary structural accounts require considerable efforts to develop. It is widely acknowledged that great artists are often perceptive and sensitive to many nuances of such complex phenomena and manage to communicate their insights without using precise structural descriptions (e.g. by relying on metaphors, associations, caricatures, dramatic narratives, poetic language, etc.). Thus, figuratively speaking, the cutting edge of art, although dull in comparison, is often ahead of the cutting edge of science, and the range of our vague artistic insights may reach some complex and nuanced psychological and social phenomena whose precise scientific descriptions may take decades to formulate if ever developed.

The above view goes against the spirit of *quietism*. Instead of declaring that ‘whereof one cannot speak, thereof one must be silent’ (Wittgenstein, 1922), the view implies that whereof one cannot speak *precisely*, one may still have something informative to say. The above admission does not undermine the epistemic authority of science. We may still hold that if an empirically plausible and scientifically comprehensive structural account is available, then it should be always preferred to its vague and imprecise alternatives. Science and art (as well as some forms of philosophy and so-called traditional wisdom) may have each their unique epistemic merits (roughly, the lesser the precision the wider the possible scope), but in the areas where their epistemic access overlaps, precise scientific accounts should be given an ultimate epistemic authority — just as precise scientific formulations are given authority over corresponding vague pre-scientific notions.

Setting epistemological questions aside, the perspectivist interpretation of non-structural qualities may also shed light on some very general moral and existential issues. As noted earlier, perspectival phenomena (both structural and non-structural) can be observed even without knowing their perspective-independent bases. But even if their perspective-independent bases are known, we may still have different attitudes towards the perspectival phenomena and their bases. For instance, we may *like* a perspectival artefact more than its

perspective-independent basis (think of the artistic impressions we choose to leave unanalysed).

In fact, in many cases the things we value most cannot be found in the perspective-independent domain — we value them *solely* for how they appear from a certain perspective — and the same is true about things we find particularly dreadful. We suffer not because we dislike the structural description of tissue damage or its behavioural manifestations, but because we dislike the perspective-dependent non-structural character of pain — and that is the primary reason why inflicting pain on others is considered morally wrong. Humphrey (2006) has suggested that the reason why sensory qualities seem so important and valuable is that their very biological function is to *matter* to us. But even if he is right, and even if we managed to spell out such a functional account in precise structural terms, it would still remain the case that the qualities *we* value are not fully structural from *our* perspective.

While there is nothing new in the idea that all values (or rather, the things we value) are to some extent perspective-dependent (they must matter *to us*), what goes usually unnoticed is that when we value feelings, sensations, impressions, and the like (or when we deem them dreadful), we do not value primarily their perspective-independent bases (even if these are known) but their perspective-dependent content that is structurally unanalysable and reductively unexplainable. So, in a sense, the perspectivist interpretation of qualities supports the popular and romantic folk-philosophical notion that there are things in the world with an utmost significance and importance that science cannot explain. Placing those things into the domain of perspectival phenomena does not reduce their value or the mystery of their existence, but it helps us understand why they can and should be disregarded in the context of fundamental physics and other natural sciences. In sum, it may be inevitable that the maximally precise and scientifically accurate theories will always leave out the most valuable and humanly important aspect of consciousness. But that is not a shortcoming of these theories. Rather, it is a consequence and manifestation of a strategic choice that has played a pivotal role in the remarkable scientific progress we have witnessed since the dawn of the modern era: the choice of focusing on experimentally accessible relational structures and dismissing the perspective-induced non-structural qualities as mere seemings.

### *Acknowledgments*

This research has been funded by the Finnish Cultural Foundation.

### **References**

- Atmanspacher, H. (2020) The Pauli–Jung conjecture and its relatives: A formally augmented outline, *Open Philosophy*, **3** (1), pp. 527–549, doi: 10.1515/opphil-2020-0138
- Beaton, M. (2013) Phenomenology and embodied action, *Constructivist Foundations*, **8** (3), pp. 298–313.
- Bohm, D. (1989) Meaning and information, in Pylkkänen, P. (ed.) *The Search for Meaning*, Wellingborough: Crucible.
- Bohm, D. & Hiley, B.J. (1993) *The Undivided Universe: An Ontological Interpretation of Quantum Theory*, London: Routledge.
- Chalmers, D.J. (1995) Facing up to the problem of consciousness, *Journal of Consciousness Studies*, **2** (3), pp. 200–219.
- Chalmers, D.J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.
- Chalmers, D.J. (2003) Consciousness and its place in nature, in Stich, S. & Warfield, F. (eds.) *The Blackwell Guide to Philosophy of Mind*, pp. 247–272, Oxford: Blackwell.
- Clark, A. (2013) The many faces of precision, *Frontiers in Psychology*, **4**, art. 270.
- Clark, A., Friston, K. & Wilkinson, S. (2019) Bayesing qualia: Consciousness as inference, not raw datum, *Journal of Consciousness Studies*, **26** (9–10), pp. 19–33.
- Crick, F. & Koch, C. (1998) Consciousness and neuroscience, *Cerebral Cortex*, **8**, pp. 97–107.
- Crick, F. & Koch, C. (2003) A framework for consciousness, *Nature Neuroscience*, **6**, pp. 119–126.
- Dennett, D. (1988) Quining qualia, in Marcel, A. & Bisiach, E. (eds.) *Consciousness in Modern Science*, pp. 42–77, Oxford: Oxford University Press.
- Dennett, D. (1991) *Consciousness Explained*, Boston, MA: Little, Brown.
- Dennett, D. (2015) Why and how does consciousness seem the way it seems?, in Metzinger, T. & Windt J.M. (eds.) *Open MIND*, **10** (T), Frankfurt am Main: MIND Group. doi: 10.15502/9783958570245
- Friston, K.J., Wiese, W. & Hobson, J.A. (2020) Sentience and the origins of consciousness: From cartesian duality to Markovian monism, *Entropy*, **22** (5), art. 516.
- Hameroff, S. & Penrose, R. (2014) Consciousness in the universe: A review of the ‘Orch OR’ theory, *Physics of Life Reviews*, **11**, pp. 39–78.
- Hautamäki, A. (2020) *Viewpoint Relativism*, Cham: Springer.
- Humphrey, N. (2006) *Seeing Red*, Cambridge, MA: Harvard University Press.
- Huron, D. (2006) *Sweet Anticipation: Music and the Psychology of Expectation*, Cambridge, MA: MIT Press.
- Koch, C. (2004) *The Quest for Consciousness: A Neurobiological Approach*, Englewood, CO: Roberts and Co.
- Ladyman, J. & Ross, D. (with Spurrett, D. & Collier, J.) (2007) *Every Thing Must Go: Metaphysics Naturalized*, Oxford: Oxford University Press.

- Loorits, K. (2014) Structural qualia: A solution to the hard problem of consciousness, *Frontiers in Psychology*, **5**, art. 237. doi: 10.3389/fpsyg.2014.00237
- Loorits, K. (2019) *Phenomenal Consciousness as a Structure in Our Brains*, Doctoral dissertation, University of Helsinki, [Online], <https://helda.helsinki.fi/handle/10138/292449>
- Markkula, G. (2015) Answering questions about consciousness by modeling perception as covert behavior, *Frontiers in Psychology*, **6**, pp. 803–815.
- Menzies, P. (2007) Causation in context, in Price, H. & Corry R. (eds.) *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, pp. 191–223, Oxford: Oxford University Press.
- Nagel, T. (1974) What is it like to be a bat?, *The Philosophical Review*, **83**, pp. 435–450.
- Oizumi, M., Albantakis, L. & Tononi, G. (2014) From the phenomenology to the mechanisms of consciousness: Integrated information theory 3.0, *PLoS Computational Biology*, **10** (5), e1003588.
- O'Regan, K. (2011) *Why Red Doesn't Sound Like a Bell: Explaining the Feel of Consciousness*, New York: Oxford University Press.
- Pylkkänen, P. (2018) Quantum theories of consciousness, in Gennaro, R.J. (ed.) *The Routledge Handbook of Consciousness*, pp. 216–231, Abingdon: Routledge.
- Robinson, W.S. (2019) *Epiphenomenal Mind: An Integrated Outlook on Sensations, Beliefs, and Pleasure*, New York: Routledge.
- Russell, B. (1927) *The Analysis of Matter*, London: Kegan Paul. Reprinted (1954), London: Allen and Unwin.
- Seager, W. (2012) Emergentist panpsychism, *Journal of Consciousness Studies*, **19** (9), pp. 19–39.
- Seth, A. (2021) The real problem(s) with panpsychism, *Journal of Consciousness Studies*, **28** (9–10), pp. 52–64.
- Strawson, G. (2006) Realistic monism: Why physicalism entails panpsychism, *Journal of Consciousness Studies*, **13** (10–11), pp. 3–31.
- Tononi, G. & Koch, C. (2015) Consciousness: Here, there and everywhere?, *Philosophical Transactions of the Royal Society B*, **370**, 20140167.
- Tononi, G., Boly, M., Massimini, M. & Koch, C. (2016) Integrated information theory: From consciousness to its physical substrate, *Nature Reviews Neuroscience*, **17**, pp. 450–461.
- Velmans, M. (2009) *Understanding Consciousness*, 2nd ed., London: Routledge.
- Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*, Ogden, C.K. (trans.), London: Routledge & Kegan Paul.

Paper received November 2021; revised May 2022.