

## The Epistemic Relevance of Morphological Content

Terry Horgan and Matjaž Potrč

Morphological content is information that is implicitly embodied in the standing structure of a cognitive system (the system's morphology) and is automatically accommodated during cognitive processing without first becoming explicit in consciousness. We maintain that much belief-formation in human cognition is *essentially morphological*: i.e., it draws heavily on large amounts of morphological content, and must do so in order to tractably accommodate the holistic evidential relevance of background information possessed by the cognitive agent. We also advocate a form of experiential evidentialism concerning epistemic justification—roughly, the view that the justification-status of an agent's beliefs is fully determined by the character of the agent's conscious experience. We two authors are each on record as defending both the thesis that much belief-formation is essentially morphological (e.g., Horgan and Tienson 1996, Henderson and Horgan 2000, Horgan and Potrč 2006, Horgan and Timmons 2007, Henderson and Horgan forthcoming), and also a version of evidentialism (Henderson, Horgan, and Potrč 2007).

Here we explain how experiential evidentialism can be smoothly and plausibly combined with the thesis that much of the cognitive processing that generates justified beliefs is essentially morphological. The leading idea is this: even though epistemically relevant morphological content does not become explicit in consciousness during the process of belief-generation, nevertheless such content does affect the overall character of conscious experience in an epistemically significant way: it is *implicit* in conscious experience, and is implicitly appreciated by the experiencing agent.

### 1. Background to the problem

In order to motivate the problem we seek to address, we will briefly sketch some reasons to think that the human cognitive processes that generate justified beliefs frequently are, and frequently must be, morphological. And we will then briefly sketch some reasons to embrace experiential evidentialism about epistemic justification.

1.1. *The holism of non-demonstrative evidence and the frame problem in cognitive science: morphological content to the rescue.*

A persistent problem in cognitive science—sometimes called the frame problem (in a broad sense of that rubric), sometimes called the relevance problem—provides strong reason to believe that the human cognitive processes that produce and sustain justified beliefs frequently make heavy use of morphological content—and that belief-fixation processes typically *must* do so, in order to be implementable in human “neural hardware.” We will briefly sketch the relevant considerations, drawing upon an important, and still very influential, discussion of the matter in Fodor (1983). (For more extensive presentations see, in addition to Fodor 1983, Horgan and Tienson 1996, Henderson and Horgan 2000, forthcoming, Fodor 2001, Horgan and Potrč 2008, and Horgan in press).<sup>1</sup>

In the closing pages of Fodor (1983), it is argued that certain problems in classical cognitive science look to be in-principle problems, and hence that the prospects for understanding processes like belief fixation within the framework of classical cognitive science are very bleak. These problems continue to plague the computational approach to the mind, and suggest the need for a radically different approach.

The main claim of Fodor’s influential book is that the human cognitive system possesses a number of important subsystems that are *modular*: domain specific, mandatory, limited in their access to other parts of the larger cognitive system, fast, and informationally encapsulated. There is good evidence, Fodor argues, that human input systems exhibit modularity. Where the classical computational approach has gotten somewhere, he says, is in understanding such modular subsystems, which by their nature delimit the class of relevant information.

Classical computationalism has made very little progress in understanding *central* processes, however. Belief fixation—the generation of new beliefs (and revising of preexisting beliefs) on the basis of current input together with other beliefs—is a paradigmatic example. Fodor argues convincingly that these processes need to have access to a wide range of cognitive subsystems, and to information on an indefinitely wide range of topics. And these same considerations, he maintains, also constitute grounds for extreme pessimism about the prospects of explaining central processes within the framework of classical computational cognitive science.

Fodor articulates these considerations in terms of the analogy between belief fixation in human cognition and scientific confirmation. Concerning central cognitive processes like belief fixation, he says, “it seems reasonable enough that something can be inferred about them from what we know about *explicit* processes of nondemonstrative inference--viz., what we know about empirical inference in science” (104). Scientific confirmation, “the nondemonstrative fixation of belief in science,” has two crucial features. It is (in Fodor's terminology) *isotropic* and *Quineian*:

By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe.... (1983, p. 105)

By saying that scientific confirmation is Quineian, I mean that the degree of confirmation assigned to any given hypothesis is sensitive to properties of the entire belief system; as it were, the shape of our whole science bears on the epistemic status of each scientific hypothesis (1983, p. 107).

Isotropy brings in the whole current belief system: any bit of actual or potential information from any portion of the belief system might, in some circumstances, be evidentially relevant to any other. Being Quineian makes confirmation holistic in a deeper way: confirmation depends upon "such considerations as simplicity, plausibility, and conservatism" (Fodor, 1983, p. 108), which are determined by the global *structure* of the whole of the current belief system and of potential successor systems.

Since belief fixation in human cognition is commonly a matter of inductive inference from the information provided by input systems and the information in memory, evidently it too must be isotropic and Quineian. Fodor concludes that it must be non-modular. He also stresses that these global aspects of belief fixation look to be at the very heart of the problems that classicism has encountered in attempting to understand such central processes:

The difficulties we encounter when we try to construct theories of central processes are just the sort we would expect to encounter if such processes are, in essential respects, Quineian/isotropic.... The crux in the construction of such theories is that there seems to be no way to delimit the sorts of informational resources which may affect, or be affected by, central

processes of problem-solving. We can't, that is to say, plausibly view the fixation of belief as effected by computations over bounded, local information structures. A graphic example of this sort of difficulty arises in AI, where it has come to be known as the "frame problem" (i.e., the problem of putting a "frame" around the set of beliefs that may need to be revised in light of specified newly available information) (Fodor, 1983, pp. 112-3).

When one considers the sorry history of attempts in philosophy of science to construct a theory of confirmation, the prospects for understanding central processing within the classical computational paradigm look very discouraging indeed. Concerning both confirmation theory and the computational conception of belief fixation, Fodor says:

The problem in both cases is to get the structure of the entire belief system to bear on individual occasions of belief fixation. We have, to put it bluntly, no computational formalisms that show us how to do this, and we have no idea how such formalisms might be developed (Fodor, 1983, pp. 128-9).

These last remarks of Fodor's remain as true today as they were in 1983, in our view. And there are two important morals to be drawn. First, central processes like belief fixation are typically too subtle and too holistic to conform to tractably implementable, exceptionless, rules for manipulating explicit representations of task-relevant information—rules of the kind that could constitute a computer program. Thus, human such cognitive processes are too subtle and too holistic to constitute, or to be modelable by, *computation* on task-relevant representations. (We ourselves think that there is at least one noncomputational framework for cognitive science that looks promising as a way of coming to terms with the frame problem—viz., that framework described by Horgan and Tienson 1996, which they call *noncomputational dynamical cognition*.)

The second moral—more directly pertinent to the present paper—is that belief fixation and related cognitive processes typically operate, and typically must operate, in a way that accommodates much relevant information *automatically* and *implicitly*. The holistic aspects of belief fixation involve not the finding and fetching of relevant representations from memory-banks where they are stored in explicit form (a method that runs afoul of isotropy), and not the overt representation and comparative evaluation of large-scale alternative belief-systems (a method that runs afoul of the Quinean dimension). Rather, these holistic aspects are somehow implicit in the structure of the cognitive system, in such a way that

temporal transitions from one occurrent total cognitive state to another accommodate the holistic aspects automatically. I.e., human belief fixation deploys—and must deploy, in order to avoid intractability—morphological processing. (Horgan and Tienson 1996 describe a way of construing morphological content, within their proposed “noncomputational dynamical cognition” framework for cognitive science. This construal involves the high-dimensional topography of *dynamical systems*—mathematical objects that can characterize, for instance, the overall potential-temporal-evolution profiles of neural networks.)

So the second moral of the frame problem is that, in general, human belief fixation *must* operate in a way that draws heavily upon morphological content, in order to avoid intractability. As we are putting it, these processes are *essentially* morphological. Belief fixation is not accomplished by manipulating explicit, occurrent, representations of all relevant information. Nor is it accomplished by mere shorthand alternatives for processes of the latter kind—say, shorthand procedures that deploy numerous special-purpose rules specifically tailored to implicitly accommodate certain background information, rather than deploying a smaller body of general-purpose rules that are applied to explicit representations of all relevant information.. (This route would encounter the frame problem all over again, now in the form of an intractably gargantuan number of special-purpose rules.) Instead, essentially morphological processing is a fundamentally different way of accommodating the holistic aspects of belief fixation.

The second moral of the frame problem is directly applicable to the notion of epistemic justification. In humans, numerous epistemically justified beliefs will have a justificatory basis that includes—and must include—morphological content. Such content will be part of the overall body of information, possessed by the agent, that constitutes *adequate evidence* for the belief. And morphological content will figure in the etiology of the belief too—albeit not by becoming explicitly present in consciousness during belief fixation, but instead by being accommodated automatically and implicitly.<sup>2</sup>

### 1.2. *Experiential evidentialism*

Evidentialism, as an approach to epistemic justification, focuses primarily on what it is to *possess* justification for a potential belief that *p* (whether or not one actually does believe that *p*)—so-called “propositional” justification. As the label suggests, evidentialism construes propositional justification as a matter of possessing *sufficiently good evidence* for proposition *p*. Evidentialists also hold that

propositional justification is necessary, although not sufficient, for one's *being justified* in believing  $p$ —so-called “doxastic” justification. Roughly and generically, the idea is that one is justified in believing  $p$  just in case (1)  $p$  is propositionally justified, relative to one's total available evidence, and (2) one believes  $p$  because of the evidential support one possesses for  $p$ .

Experiential evidentialism, as a species of generic evidentialism, holds that the degree of propositional justification an agent possesses for a proposition  $p$  ultimately is fully determined by the agent's *experience*. The leading idea, roughly and generically formulated, is this:  $p$  is propositionally justified, for an agent  $A$  at time  $t$ , just in case  $p$  is *very likely true, given  $A$ 's experiential situation at  $t$* . Different forms of experiential evidentialism will cash out this generic characterization in different ways—depending, for instance, on (i) whether or not one counts  $A$ 's prior experiential history (or portions thereof) as part of one's experiential situation *at  $t$* , (ii) whether or not one treats mere dispositions (at  $t$ ) toward certain kinds of experiences as part of one's experiential situation at  $t$ , (iii) how one construes the idea that propositional justification is fully determined by experience, (iv) whether or not one builds in foundationalism about the structure of propositional justification (with foundational beliefs being ones that bear some specially intimate connection to experience—e.g., beliefs about the nature or content of experiences, of beliefs whose content matches the content of experiences themselves), and (v) how one construes the rubric *experience*.

Although the problem we will pose and address in this paper probably can arise in connection with most any version of experiential evidentialism, we will focus here on a specific version that we ourselves find especially plausible—and for which the problem will arise in a very vivid way. This version makes the following claims concerning the lately-mentioned items (i)-(iv).

Regarding (i): Agent  $A$ 's experiential situation at time  $t$  is entirely *synchronic*, involving only how things are with  $A$  *at  $t$* . Earlier experiences are only epistemically relevant, at  $t$ , insofar as they are somehow reflected in how things are synchronically with  $A$  at  $t$ —e.g., in the form of synchronic memory-experiences of those earlier experiences. This synchronic version seems to us more in the spirit of experiential evidentialism than would be a version that allows, for instance, past experiences to count as part of one's present experiential situation even if those experiences are not in any way reflected in the overall character of one's present experience.

Regarding (ii): Agent A's experiential situation at time  $t$  includes not only A's total occurrent experiential state at  $t$ , but also A's synchronic dispositions, at  $t$ , to undergo *potential* total occurrent experiential states triggered by *potential* experiential inputs (e.g., triggered by potential sensory-perceptual inputs). This stipulation makes for a version of experiential evidentialism that covers not only A's occurrent beliefs at  $t$ , but also A's beliefs at  $t$  that are merely dispositional rather than occurrent. If A were to undergo experiences that triggered an occurrent token (at time  $t + \delta$ ) of A's dispositional belief that  $p$ , then A's justification for that token belief would be entirely determined by A's total synchronic experiential state (at  $t + \delta$ ).

Regarding (iii): The idea that agent A's propositional justification for  $p$  is fully determined by A's (actual or dispositional) synchronic experience is to be understood as a supervenience claim: for any possible agent A\* who occurrently or dispositionally believes  $p$  (at time  $t^*$ ) and who is in the same experiential situation (at  $t^*$ ) as is A (at  $t$ ), A\* (at  $t^*$ ) has the same propositional justification for  $p$  as does A (at  $t$ ). In short: no difference in justification without a difference in synchronic experience.

Regarding (iv): We do *not* stipulate that experiential evidentialism is to be a version of foundationalism about the structure of justification. This is because we believe that there is a viable and attractive version of experiential evidentialism—one incorporating the claims lately made regarding issues (i)-(iii) and also the remarks just below about item (v)—that embraces a form of coherentism about the structure of epistemic justification, rather than embracing foundationalism. (We return to this theme in section 4 below.)

Regarding (v): The words 'experience' and 'experiential', as here employed, are to be understood non-restrictively, as referring to those aspects of occurrent mentality that are conscious in a broad sense—i.e., conscious *as opposed to unconscious*. This leaves open a host of important questions in philosophy of mind—e.g., whether *phenomenal* consciousness coincides with consciousness in this broad sense or instead is only a proper part of it, whether occurrent propositional attitudes have distinctive phenomenal character, and whether sensory phenomenal experience has intentional content. For present purposes, we can remain neutral about such questions, although we ourselves are inclined to be very liberal about the existence of non-sensory phenomenology (including cognitive phenomenology and agentic phenomenology) and about the inherent intentionality of phenomenology. Some versions of experiential evidentialism—in particular, some foundationalist versions—construe the rubric *experience* more

narrowly, e.g., as comprising only sensory phenomenology (perhaps regarded as having intentional content, or perhaps not). We adopt a more inclusive construal here for two reasons, among others: first, our lately-mentioned inclination toward liberalism about the extent of phenomenal consciousness, and second, our recently-mentioned inclination toward a non-foundationalist form of experiential evidentialism (a form that has no need to delineate a class of narrowly construed “experiences” that figure in the subject matter of putatively foundational beliefs).

Having set out the form of experiential evidentialism we will focus upon here, hereafter we will mostly (except in section 5) use the expression ‘experiential evidentialism’ to refer to this specific version of the generic view. Let us now briefly mention some considerations in favor of the position. (The remarks to follow are meant to quickly provide some motivation, not to constitute a rigorous and detailed defense. Also, they do not necessarily provide direct support for the specific version under discussion, as opposed to others.)

First, consider the familiar thought-experimental scenario of an envatted brain hooked to a computer that generates its sensory inputs, monitors its motor outputs, and does so in such a way that brain’s neural activity throughout its lifetime exactly matches that of an ordinary person’s brain throughout that person’s life. Many people, ourselves included, have very strong pre-theoretic intuitions about this scenario, concerning both the character of the envatted brain’s mental life and the epistemic status of its beliefs.

Regarding mentality, the intuition is that there is a very strong match between the envatted brain’s mental life and that of the ordinary person—including a strong match in the beliefs possessed, and in the intentional contents of those beliefs. The beliefs have matching “narrow content,” even though at least some of them (e.g., the ones that purport to refer to individuals, and/or to natural kinds) might differ with respect to “wide content.” (Plausibly, the pertinent beliefs in the envatted brain do not successfully refer to any individuals or natural kinds at all, and thus lack wide content altogether.) This goes contrary to recently-popular versions of strong externalism about mental intentionality, but so much the worse for those positions (cf. e.g. Horgan and Tienson 2002, Horgan, Tienson and Graham 2004).

Regarding justification, the intuition is that the matching beliefs of the envatted brain and the ordinary person have exactly the same justification. Not only do they have the same *degree* of justification, but they also have the *same* justification—by virtue of the fact that the envatted brain and

the ordinary human have exactly similar experience (both occurrently and dispositionally). And herein lies one, quite powerful, motivation for experiential evidentialism: this position conforms with the strong pre-theoretic intuition that the envatted brain's beliefs have the same justification as do the corresponding beliefs of the normal human, by virtue of the fact that both of these creatures have exactly similar experiences and experience-dispositions.

A second consideration favoring experiential evidentialism is the fact that the normative notion of epistemic justification is very often deployed in a *deontological* way, i.e., a way that is intimately connected to the notion of responsibility.<sup>3</sup> The idea is that justified beliefs are those that are *responsibly held, given the agent's available evidence*—whereas unjustified beliefs are those that are held irresponsibly. This idea leads directly to the judgment that the corresponding beliefs of the envatted brain and the normal human have exactly matching degrees of epistemic justification—since an epistemic agent's degree of responsibility, in holding a given belief, is surely dependent only upon matters within the purview of the agent's experience. In short: matching experience makes for matching degrees of epistemic responsibility, hence matching degrees of justification. Experiential evidentialism honors this intuition.

## **2. The problem: How can experiential evidentialism recognize the epistemic relevance of morphological content?**

We have argued that much human belief-formation must draw heavily on large amounts of morphological content, in order to tractably accommodate the holistic evidential relevance of background information available to the agent. The pertinent background information is part of the agent's *evidence*; hence, presumably it should count not merely as causally operative in the etiology of the belief, but also as epistemically relevant to the justificatory status of the belief. We will call this contention *the epistemic relevance of morphological content*.

We have also argued in favor of experiential evidentialism, which asserts that the justificatory status of an agent's belief—comprising both its degree of justification and the specific evidential support it possesses—is supervenient on the agent's occurrent synchronic experience. On this account of epistemic justification, any psychological factors that figure causally in the generation of a belief that is occurrent at

time  $t$ , but are not part of the agent's synchronic experience at  $t$ , are irrelevant to the justificatory status (at  $t$ ) of that occurrent belief.

A problem now arises: *prima facie*, experiential evidentialism looks to be incompatible with the epistemic relevance of morphological content. The apparent incompatibility arises in the wake of the following two additional claims, each of looks *prima facie* plausible itself. (1) If an epistemic agent has occurrently believes that  $p$  at time  $t$ , and an item of background information figures in the formation of this occurrent belief only morphologically rather than by becoming consciously explicit, then role of that item of morphological content in the etiology of belief is independent of the synchronic character (at  $t$ ) of the agent's experience. (2) Under experiential evidentialism, information in the cognitive system that is independent of the agent's synchronic experience (at a time  $t$ ) is not epistemically relevant to any occurrent belief (at  $t$ ).

We will propose a solution to this problem, involving two keys ideas. First, deny claim (1). Second, claim instead that epistemically relevant information, and appreciation of the epistemic relevance of such information, can figure as aspects of synchronic conscious experience without being consciously explicit.

### **3. The solution: Chromatic illumination as the implicit presence in consciousness of morphological content.**

A visual metaphor will be useful as a way of articulating our proposal. Consider a visual scene that is illuminated in certain ways by light sources that are not themselves visible (from the observer's perspective) within the scene, and that significantly affect the overall look of the scene.

Think, for instance, of the famous 1892 painting by Toulouse Lautrec, "At the Moulin Rouge," which hangs in the Art Institute of Chicago. Various figures in the painted scene are illuminated in strikingly different ways—e.g., the women more prominently than the men, one woman by lighting to the left but outside the scene, another woman by lighting from the lower right but outside the scene, a peculiar light-induced greenish tint to some of the illuminated faces that blends with the greens in the background of the scene, etc. The presumptive sources of these distinctive features—lighting of various kinds at various positions in the presumptive wider environment, producing light with various different chromatic characteristics—are not present in the visible scene. They are not explicitly represented.

Nonetheless, they are *implicitly* present in the scene anyway, in the ways that the figures in the scene are chromatically illuminated by those presumptive light-sources.

The directly visible scene presented in a painting—or in a photograph, or on a stage—can be taken as a metaphorical stand-in for what is explicitly present in conscious experience. By contrast, the out-of-view sources of the visible scene’s various aspects of chromatic illumination can be taken as metaphorical stand-in for what is implicitly present in consciousness—present not by being explicitly presented or depicted, but rather by virtue of how it affects the character of what *is* there explicitly. These effects, within synchronic consciousness, of information not explicitly present in consciousness, are what we will call *chromatic illumination*.

### 3.1. *Joke-getting as a philosophical paradigm*

Before turning directly to the matter at hand, it will be useful to describe a familiar kind of experience that plausibly exhibits the phenomenon of chromatic illumination—and that involves a form of mental processing that is not essentially tied to vision, or to any particular sense modality. (The advantage of focusing on a non-sensory cognitive phenomenon is that it is appropriately analogous, so we will maintain, to belief formation.)

Consider a familiar type of experience: understanding a joke. For specificity, consider a cartoon by Danny Shanahan that appeared in the *New Yorker* magazine roughly a dozen years ago. (We will describe it rather than exhibiting it—which underscores the fact that pertinent aspects of joke-getting normally do not depend essentially on specific sense modalities.) Two guys are sitting in a bar, and the one who is speaking has an unhappy and exasperated expression on his face. He says to the other guy, “What I don’t, like, get is how she, like, figured out I was, like, having an affair with the, like, babysitter.”

Getting this joke, or any joke, is an instantaneous experience. In that instant, normally one needs to appreciate quite a wide range of pertinent background information; also, one needs to appreciate why and how all this information combines, holistically, to constitute an instance of funniness. Each item of that information must be appreciated, together with the ways that all the pertinent items interact with one another to make for funniness—for, otherwise one would not be understanding the joke.

For any reasonably clever joke, typically it is possible to elaborate at some length upon the various items of background information all of which are pertinent. In the example at hand, for instance,

probably the most salient such item—the item that one would mention first, in seeking to explain the joke to someone who does not yet get it—is that the practice of persistently inserting the word ‘like’ into one’s spoken English is a distinctive feature of way *teenage girls* talked in the U.S., a dozen years ago.

(Supposedly the practice was initiated by teenage females in the San Fernando Valley, adjacent to Los Angeles—so-called “valley girls,” who are sources of numerous youthful fads in the U.S. Nowadays, such ‘like’-talk is much more prevalent: it is the norm among college students in the U.S. when they are speaking with peers or with intimates, and it is relatively common among young adults.)

But that is just the beginning, in terms of pertinent background information that must be grasped in the instant in order to get the joke. Additional such items of information can be made salient by posing suitable questions—questions that might be asked by someone who is persistently deficient at understanding jokes. Whose children did this babysitter care for? Roughly how hold is the babysitter? How did the opportunity arise for the guy to have an affair with the babysitter? Who is the “she” who figured out that the guy is having an affair with the babysitter? How did that person react, upon learning this fact? Why is the guy upset that she figured it out and reacted that way? Why doesn’t he himself understand how she figured it out? And, of course: What’s so funny about all this? One could write a monograph about the pertinent background information and its holistic relevance. Likewise for most any even moderately subtle joke.

In the instant of joke-getting, very little of this kind of information seems to be explicitly present in consciousness. Nonetheless, all of it is being *appreciated* in one’s conscious experience, because otherwise one would not get the joke. Thus, the information is *implicitly* present in consciousness, by way of the chromatic illumination of one’s overall synchronic experience. Furthermore, the “getting it” aspect of experience is not some generic feature, such as experiencing oneself laughing or inclined to laugh (perhaps without knowing why). Rather, it is quite content-specific: some particular item(s) of explicitly conscious content (in this case, what the guy at the bar is saying, and his obvious consternation) is appreciated as funny by virtue of how those explicitly conscious items are relevantly interconnected with a rich body of pertinent background information. Thus, all those *specific* items of background information are implicitly present in the conscious joke-getting experience, by virtue of the *specific* way that the experience is chromatically illuminated.

Now admittedly, it is an empirical hypothesis that joke-getting has the features lately described. But it is a very plausible hypothesis, for at least two mutually supporting reasons. First, it conforms with the deliverances of introspection: when one attends introspectively to one's joke-getting experience (albeit perhaps retrospectively via one's memory of the recent past), it does not seem that all that background information is explicitly present in consciousness; yet, it also seems, upon reflection, that one would not have understood the joke without consciously *appreciating* all that information and its holistic humor-relevance.<sup>4</sup> Second, there are strong theoretical grounds—really the same grounds as in the case of belief-fixation—for the claim that the cognitive processes that underwrite joke-getting are essentially morphological: viz., joke-getting exhibits very similar kinds of holistic, Quineian-cum-isotropic, dependence on background information. This means that it is not tractably possible for joke-getting cognition to proceed by deploying explicit conscious representations of all pertinent items of information and their holistic interconnections—especially not in the very short time it normally takes to get a joke.

We will use the expression 'implicit conscious content' for information that is implicitly present in synchronic experience as synchronic illumination. This notion deserves elucidation and elaboration, ideally from multiple perspectives: introspective phenomenological description, philosophical conceptual analysis, cognitive-scientific inquiry both theoretical and empirical. And it has potential relevance to a variety of topics and issues in philosophy of mind, in cognitive science, and in epistemology. It deserves serious exploration, beyond what we can undertake here. However, we will offer some tentative remarks about it, before proceeding to the business at hand in the present paper.

What can one say, by way of introspective phenomenological reflection, about the nature of chromatic illumination? For specificity, let us continue to focus on the experience of getting the joke in Shanahan's 'like'-talk cartoon. Two claims about the chromatic-illumination character of this experience both seem very plausible, from the introspective perspective. First, there could not be a creature whose synchronic experience is exactly like one's own but who nevertheless *failed* to get the joke; i.e., really getting it is a feature that is supervenient on one's overall synchronic experience. Such supervenience gives substance to the claim that the pertinent background information is *implicitly present* in one's synchronic experience: there could not be a creature who has this same experience but fails to appreciate that information.<sup>5</sup>

Second, one's appreciation of implicit background information manifests itself experientially, in part, via an aspect of "looming potentialities" concerning the joke, viz., a sense that one could, if suitably prompted by others or by oneself, manifest one's appreciation of such information—in overt linguistic behavior, and/or by explicitly bringing such information to mind in one's conscious thought. One kind of looming potentiality, for example, is the capacity to spontaneously arrive at consciously explicit answers to suitable background-probing questions. (Roughly how old is the babysitter? Who is the "she" who figured out that the guy is having the affair with babysitter? Etc.) Also, the presence of these looming potentialities is normally a matter of chromatic illumination itself, rather than something that is explicitly represented in consciousness.

Must an experiencer *actually* possess such understanding-manifestation capacities, in order to really understand the joke? This is a subtle and delicate matter. On one hand, there is a strong inclination to say that the possession of suitable dispositions is partly constitutive of genuine understanding (as opposed to mere seeming-to-understand). On the other hand, presumably there are nomically possible creatures who are synchronic-experience duplicates of an ordinary human joke-getter but are quite lacking in such dispositions—e.g., a partial-brain, envatted and only momentarily conscious, that lacks not only a body but also those parts of the brain that are the categorical basis for the relevant behavioral and cognitive dispositions. What is right to say about such a creature, we suggest, is that it has suitable *conditional* understanding-manifesting dispositions—which means, roughly, that if the creature were the kind of being that its experience overtly represents it (to itself) to be, and were embodied and en-worlded in the way that its experience overtly represents it to be, then it would possess appropriate understanding-manifestation dispositions. The claim, then, is this: possession of conditional dispositions toward suitable understanding-manifestation is constitutive of genuine understanding (as opposed to mere seeming-to-understand), and furthermore is supervenient on the joke-getter's synchronic experience. It is a constitutive aspect of the experience's chromatic illumination.

Are there determinately *specific* understanding-manifesting conditional dispositions that the experiencer must possess in order to have genuine understanding, and that supervene on synchronic experience? We find it plausible, on the basis of reflective introspection, that the answer to this question is no. What's supervenient on synchronic experience is that *there are* suitable understanding-manifesting conditional dispositions that the experiencer possesses. But just what those dispositions are depends in

part on aspects of the agent's cognitive architecture and bodily endowments that are not necessarily included, even implicitly, in the content of synchronic experience.<sup>6</sup>

The remarks in the preceding four paragraphs are a brief attempt to say something more specific and clarificatory, in order to further elucidate the kind of chromatic illumination that figures in experiences like joke-getting. What we will say below does not require these somewhat tentative remarks to be fully correct, however; the notion of chromatic illumination can serve our present purposes anyway. The nature of chromatic illumination, and of the implicit content that supervenes on it, is a topic well worth further exploration.

### 3.2. *Chromatic illumination and epistemically relevant morphological content.*

We are ready now to address the conundrum described in section 2. The problem, remember, is to reconcile experiential evidentialism with the epistemic relevance of morphological content. Our reconciliatory proposal is as follows. (We will state it for occurrent beliefs, but it extends straightforwardly to dispositional beliefs.) Suppose that an agent has an occurrent belief that  $p$ , at time  $t$ . Any morphological content that figures in the justification of this belief does more than play a causal role in generating it. In addition, such content-items, together with an appreciation of their holistic evidential import vis-à-vis the proposition  $p$ , are implicitly present in the agent's synchronic experience—present as chromatic illumination. This means, inter alia, that the same information would be implicitly present in the synchronic experience of any possible synchronic-experiential duplicate of the given agent. (Implicit content is supervenient on synchronic experience.<sup>7</sup>)

What can be said about the character of such chromatic illumination? One important feature of it, somewhat analogous to the kind of chromatic illumination that figures in joke-getting, is that it imbues one's overall experiential state with a sense of the *evidential fittingness* of the belief that  $p$ —often including an aspect of *cohering well* with other aspects of experience, both sensory-perceptual and doxastic. Moreover, what the chromatic illumination imbues is not some purely generic evidential-fittingness aspect (so that one could experience the evidential fittingness of proposition  $p$  while yet having no idea what specific evidential support there is for  $p$ ), but rather a sense of *content-specific* fittingness. This contrasts, for instance, with the quite different chromatic illumination that would accompany occurrent beliefs experienced as arising “out of the blue.”

Closely related to this sense of evidential fittingness, and again somewhat analogous to what happens in joke-getting, is the sense of looming potentialities—in this case, potentialities for exhibiting one’s appreciation of pertinent evidential support possessed by proposition  $p$ . It is probably not fully determinate just what those looming potentialities are—since this may well depend partly upon aspects of one’s cognitive architecture that are not part of either the explicit content or the implicit content of synchronic experience. The potentialities might include certain capacities for articulating aspects of justificatory support, if queried to do so (although some people are much better at this than others); in any case, they very likely include certain capacities for giving appropriate answers to specific questions that might be posed by way of probing one’s appreciation of the evidence for  $p$ .

The appeal to chromatic illumination thus solves our reconciliation problem. In addition, it also has at two other significant advantages in its favor. First, it conforms well with the deliverances of introspective phenomenological inquiry, when one directs such inquiry specifically at what it is like to undergo an occurrent belief that seems to oneself to be well justified. Second, it avoids being an overly intellectualized version of experiential evidentialism, a version that would implausibly require the entire body of information that is evidentially relevant to the belief that  $p$ , and the full Quinean/isotropic structure of the evidential relevance itself, to be explicitly and synchronically present in consciousness.

#### **4. On the structure of epistemic justification**

Traditionally, versions of experiential evidentialism usually have been wedded to foundationalism about the global structure of epistemic justification. Roughly and generically, the idea is that there is a class of epistemically basic beliefs, whose content involves synchronic experience itself; and that all other justified beliefs acquire their epistemic justification by way of inferential connections (both demonstrative and nondemonstrative) from basic beliefs to other beliefs in an agent’s doxastic system.

Although the position we have advocated here is consistent with foundationalism about the global structure of epistemic justification, it is not committed to foundationalism. The position also can be wedded instead to coherentism. We ourselves are inclined to favor a version of coherentism that puts heavy weight not only on coherence per se, but also on the (defeasible) requirement that an epistemically appropriate doxastic system should include beliefs that are very largely in accord with the contents of sensory-perceptual experience. (The requirement is defeasible because one could conceivably have good

experiential evidence that most of one's sensory-perceptual experience is illusory and nonveridical.) Conventional versions of coherentism typically are unrealistically intellectualistic, and wildly so: they require, as a prerequisite for justified belief, that the epistemic agent hold explicitly in mind the whole doxastic system and its coherentistic features. But that problem is avoided by the position we have described here. Conscious appreciation of global coherence in one's doxastic system, and conscious appreciation of how the occurrent belief that  $p$  fits coherence-wise with one's overall doxastic system, could well be implicitly present in synchronic consciousness via chromatic illumination—even though it is not possible to explicitly represent all this information at once in synchronic consciousness.<sup>8</sup> (For articulation and defense of a version of coherentism that relies heavily upon morphological content, see Henderson and Horgan 2000, forthcoming.)

### 5. On contextualist pluralism about epistemic justification.

One possibility that is consistent with all we have said here is that sometimes, some of the morphological content that is evidentially relevant to the proposition  $p$ , and that figures in the cognitive processing that causally generates the occurrent belief that  $p$ , was instilled in the standing structure of the epistemic agent's cognitive system via *prior* conscious experience that is not reflected at all in the agent's present synchronic conscious experience—not even implicitly, via chromatic illumination.

Consider, for example, a familiar kind of situation: occurrently believing that  $p$  (at time  $t$ ), where the attendant chromatic illumination is as-of *remembering* that  $p$ , but one has no inkling of how or when one acquired the information that  $p$ . Suppose that at least sometimes in such cases, (1)  $p$  originally became an item of “beliefish” morphological content<sup>9</sup> in the standing structure of one's cognitive system, by virtue of prior experiences that, when they occurred, constituted good evidential support for  $p$ , but (2) those experiences themselves have not been retained in memory, and thus (3) those experiences are no part of the overall content (either explicit content or implicit content) of one's total experiential state *at time t*.

Now, according to the synchronic version of experiential evidentialism we have been discussing here, the prior experiences that installed the beliefish morphological content  $p$  into the standing structure of the agent's cognitive system, and that constituted good evidential support for  $p$ , are irrelevant to the agent's *present* justification for occurrently believing  $p$ . Any being with exactly similar synchronic

experience would have exactly the same justification for the occurrent belief that  $p$ —even such a being who came into existence five minutes ago and never actually had any prior experiences that constituted good evidence for  $p$ . (Call a being like that a *pastless synchronic-experiential duplicate* (for short, a PSE duplicate).

Is this the right verdict, concerning the justification-status of the occurrent belief that  $p$  in the PSE duplicate? We ourselves find that our intuitions are pulled in two conflicting directions. On one hand, when deploying the notion of epistemic justification in a synchronically deontological way—with an eye toward whether or not the cognitive agent is *now* being *epistemically responsible* in occurrently believing that  $p$ —we are inclined to say yes: the PSE duplicate has the same justification for believing  $p$  as does the ordinary agent, justification that is supervenient on the overall character of synchronic experience while occurrently believing that  $p$ .

On the other hand, in some contexts of normative epistemic appraisal, it seems natural to deploy the notion of epistemic justification in a way that takes into account certain diachronic facts about the cognitive agent and about the diachronic etiological history of the agent's occurrent belief (at  $t$ ) that  $p$ . Suppose, for instance, that the beliefish morphological content that  $p$  originally became installed into the agent's cognitive system, at time  $t-\delta$ , by virtue of the agent's forming an occurrent belief that  $p$  (at  $t-\delta$ ) that was unjustified (at  $t-\delta$ ). Sometimes it is contextually appropriate to do deontic epistemic evaluation diachronically, rather than synchronically. From that perspective, the agent's present occurrent belief that  $p$  is unjustified—despite the agent's seeming to remember that  $p$ —because the beliefish morphological content that  $p$  was originally was incorporated into the agent's cognitive system as a byproduct of an earlier, epistemically irresponsible, occurrent belief that  $p$ . And sometimes it is contextually appropriate to do epistemic evaluation in a way non-deontic way that focuses primarily on the integrity of the cognitive mechanisms that generated the occurrent belief that  $p$  (at  $t$ ), rather than on whether or not the agent is responsible in holding the belief. From that perspective too, the agent's present occurrent belief that  $p$  is unjustified: the processes that generated (at  $t-\delta$ ) the earlier occurrent belief that  $p$  were deficient because unreliable and not suitably evidence-sensitive, which means that the overall diachronic history of the present occurrent belief that  $p$  is also epistemically deficient.

Can one accommodate these multiple normative perspectives, and in a way that fits smoothly with the position developed in sections 1-4? We would say yes. Let us briefly sketch a way of doing so,

one that incorporates the kind of coherentism about the structure of justification that was described in section 4. (We will focus on occurrent beliefs, but the story can be extended fairly straightforwardly to dispositional beliefs too.) Suppose an epistemic agent has an occurrent belief that  $p$ , at time  $t$ . Consider, first, the content that figures in the agent's overall synchronic experience (at  $t$ )—both explicit content and implicit content. Insofar as the pertinent items of explicit and implicit synchronic content bear suitable, holistic, coherence relations to one another—including cohering well with the contents of synchronic sensory-perceptual experience—the agent's synchronic occurrent beliefs count as epistemically justified, by the standards of experiential evidentialism. (Occurrent beliefs based on seeming-memories, for instance, would cohere well with other aspects of synchronic content, even if the agent only came into existence a moment ago. Likewise for occurrent beliefs that were produced by beliefish morphological content that itself originally got installed in the cognitive system via defective processes that were not evidence-respecting.)

Now consider a wider body of items of explicit and implicit content—those that have been synchronically present at some time or other in the agent's experiential history up to the present time  $t$ . Insofar as the pertinent items of explicit and implicit content within this diachronic experiential history bear suitable, holistic, sensory-perceptual-experience respecting, coherence relations to one another, the agent's occurrent beliefs that are present within this experiential history count as epistemically justified.

From this diachronic experiential-evidentialist vantage point, ascriptions of epistemic justification will work somewhat differently than they do under the synchronic version of experiential evidentialism we focused upon in sections 1-4. Some synchronic beliefs that count as justified under synchronic experiential evidentialism will count as unjustified under the diachronic version—e.g., a synchronic occurrent belief that is based on a synchronic seeming-memory, where one has no inkling how or when one first adopted the belief, and in fact it originally arose without suitable synchronic-experiential justification. And some beliefs that count as adequately justified, under synchronic experiential evidentialism, will count from diachronic point of view as *more* justified—e.g., a synchronic occurrent belief that is based on a synchronic seeming-memory, where one has no inkling how or when one first adopted the belief, and in fact it originally arose *with* suitable synchronic-experiential justification. (From the diachronic point of view, this last kind of agent is better justified in believing that  $p$  than is a PSE duplicate who just came into existence.)

The evaluative perspective just described is naturally viewed as still invoking a variant of experiential evidentialism, albeit a diachronic one. But in addition, there may well be contexts in which it is appropriate to assess justificatory status in a way that incorporates yet further items of content that are not supervenient on experience or experiential history. For instance, it is very plausible there is pertinent information in the overall cognitive system of a normal human that one hand is never present in consciousness (not even implicitly), and on the other hand contributes yet further evidential support to the agent's beliefs. Suppose, for example, that the early, pre-conscious stages of visual processing operate in something like the manner described in David Marr's famous computational theory of vision (Marr 1982). Some of the pertinent information might get explicitly represented in the successive pre-conscious stages, although without being conscious either explicitly or implicitly. And some of it might well be morphological, and again might not ever enter consciousness (not even implicitly).<sup>10</sup> Nonetheless, the content that figures in successive stages of early vision, in a Marr-type theory, does bear evidential relations to the conscious contents of visual experience—and thereby also bears evidential relations to occurrent visual beliefs. So for some purposes, such as inquiry into the epistemically competent workings of human perceptual and cognitive architecture, it would seem quite appropriate to treat epistemic justification as a matter of fitting well into a suitably coherent, suitably perceptual-experience respecting, overall system of intentional mental states and structures that goes beyond what is explicit or implicit in the agent's synchronic experience or diachronic experiential history. This is arguably still an evidentialist in spirit, albeit no longer a purely *experientialist* kind of evidentialism. Also, it is no longer suitable for deontic evaluation (either synchronic or diachronic)—although it perhaps still exhibits a cognitive-engineering oriented kind of deontology in which one asks whether the *cognitive system* is operating the way it “ought to operate.” And from this yet-more-inclusive point of view, an actual human will count as being better justified in many or most beliefs than is, say, a creature who has a matching experiential history but does not undergo anything like the richly intentional, deeply unconscious, evidentially pertinent cognitive processes that operate in the early stages of normal human vision. (Maybe this diachronic experiential duplicate is an envatted *partial* brain that lacks visual-input channels and instead has external stimulation fed directly into its visual cortex.)

So, in terms of accommodating the range of somewhat divergent judgments and judgment-tendencies concerning how to assess the epistemic justification of agents' beliefs in various

circumstances, and in terms of accommodating the various different purposes and perspectives that can legitimately inform the evaluative use of the notion of epistemic justification, there is good reason to embrace the kind of contextualist pluralism just sketched. This means allowing for some contextually appropriate attributions of epistemic justification that do not fully conform with the version of experiential evidentialism we espoused in earlier sections. In some contexts, the operative epistemic standards can conform instead to a diachronic version of experiential evidentialism. And in still other contexts, they can conform to standards that allow for the justificatory relevance of forms of evidence within the cognitive system that are outside the scope of conscious experience altogether.

## **6. Conclusion**

Human cognitive processes that generate justified beliefs are frequently morphological. We motivate this by citing holism of non-demonstrative evidence and the frame problem in cognitive science; arguable, these features can be tractably accommodated in belief-updating only via extensive reliance on morphological content. We also subscribe to experiential evidentialism regarding epistemic justification, motivated by the deontic element in epistemic appraisal and by intuitions that there is a match in belief-content and in justification between an ordinary human and an experiential duplicate that is an envatted brain. The problem now arises how experiential evidentialism can recognize the epistemic relevance of morphological content, given that, as it seems, if an item of information figures in occurrent belief only morphologically rather than becoming consciously explicit, it is independent of the synchronic character of the agent's experience, and thus cannot be epistemically relevant. We oppose this presupposition by introducing chromatic illumination as the implicit presence in consciousness of morphological content, illustrating it first with an example from visual arts, and then introducing joke-getting as a philosophical paradigm. Joke-getting is an instantaneous, specific experience involving a large amount of holistic background information, appreciated in one's conscious experience. Getting the joke is a feature that is supervenient upon one's overall synchronic experience. Morphological content and appreciation of its holistic evidential support are present in the agent's synchronic experience as chromatic illumination. The global structure epistemic justification, we suggest, is coherentist rather than foundationalist. Finally, we endorse contextualist pluralism about epistemic justification, thereby accommodating contextually

appropriate uses in which epistemic appraisal as diachronically deontic rather than synchronically deontic, and contextually appropriate uses in which responsible epistemic agency is not directly at issue.<sup>11</sup>

## Literature

Fodor, J. (1983). *The Modularity of Mind*. MIT Press.

Henderson, D. and Horgan, T. (2000). Iceberg Epistemology, *Philosophy and Phenomenological Research* 61, 497-535.

Henderson, D. and Horgan, T. (forthcoming). *The Epistemological Spectrum: At the Interface of Cognitive Science and Conceptual Analysis*.

Henderson, D., Horgan, T., and Potrč, M. (2007). Transglobal Evidentialism-Reliabilism, *Acta Analytica* 22, 281-300.

Horgan, T. (in press). Connectionism, Dynamical Cognition, and Non-Classical Compositional Representation, in E. Machery, W. Hinzen, and W. Werning (eds.), *The Oxford Handbook of Compositionality*. Oxford University Press.

Horgan, T. and Potrč, M. (2006). Particularist Semantic Normativity (with M. Potrc). *Acta Analytica* 21, 45-61.

Horgan, T. and Potrč, M. (2008). *Austere Realism: Contextual Semantics Meets Minimal Ontology*. MIT Press.

Horgan, T. and Timmons, M. (2007). Morphological Rationalism and the Psychology of Moral Judgment (with M. Timmons), *Ethical Theory and Moral Practice* 10, 279-295.

Horgan, T. and Tienson, J. (1996). *Connectionism and the Philosophy of Psychology*, MIT Press.

Horgan, T. and Tienson, J. (2002). The Intentionality of Phenomenology and the Phenomenology of Intentionality, in D. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press, 520-533.

Horgan, T., Tienson, J., and Graham, G. (2004). Phenomenal Intentionality an the Brain in a Vat, in R. Schantz (ed.), *The Externalist Challenge*. Walter de Gruyter, 297-317.

Potrč, M. (2000). Justification Having and Morphological Content, *Acta Analytica* 24, 151-173.

---

<sup>1</sup> Much of the material in this subsection is adapted from other texts in which it is also used, including Horgan and Tienson 1996 and Henderson and Horgan 2000.

<sup>2</sup> Does the morphological embodiment of an evidentially pertinent item of information I constitute a *belief* with content I? Often enough, we would maintain, the answer is yes. This means that the traditional dichotomy between occurrent beliefs and dispositional beliefs is not exhaustive (where a dispositional belief that p is construed, roughly, as a disposition to occurrently believe p in suitable elicitation circumstances). There are also morphological beliefs, which play a role in cognition by being automatically accommodated without becoming occurrent along the way. (Needless to say, one and the same belief *type* could get tokened, at different times in a creature's life, in any or all of the three ways—sometimes occurrently, sometimes dispositionally, and sometimes morphologically.)

<sup>3</sup> The notion of epistemic justification sometimes can be appropriately deployed for deontic normative appraisal that is diachronic rather than synchronic, and sometimes can be deployed from a perspective that does not directly involve responsible epistemic agency. We take up these matters in section 5 below.

<sup>4</sup> It is conceivable that all pertinent background information is explicitly present in consciousness, in such a way that most of it simply is not noticed or attended to. But the burden of proof is on those who would seek to defend that alternative hypothesis.

<sup>5</sup> The kind of intentional content (both explicit and implicit) that supervenient on experience is *narrow* content. For a conception of narrow content that we ourselves would favor, see for instance Horgan, Tienson, and Graham (2004).

<sup>6</sup> What about the above-mentioned case of a partial-brain, envatted and only momentarily conscious, that lacks not only a body but also those parts of the brain that are the categorical basis for the relevant behavioral and cognitive dispositions? What seems right to say about this creature, at least roughly, is this: if it were a being of the kind its experience overtly represents it to be, embodied and en-worlded in way its experience overtly represents it to be, then it would have *some* kind of cognitive-architectural scaffolding that subserves *some* kind of suitable understanding-manifesting dispositions.

<sup>7</sup> We again emphasize that kind of implicit content we are talking about is narrow; cf. note 5.

<sup>8</sup> Some or all implicitly present content presumably is subject to justificational appraisal itself, especially insofar as it qualifies as implicit belief. The coherentistic story extends to that too. Likewise, there can be epistemically *unjustified* implicit content too.

<sup>9</sup> What we are calling a “beliefish” item of morphological content is something that we ourselves would consider a morphological *belief*; cf. note 1. But for present purposes it doesn’t matter whether or not it counts as a full-fledged belief, which is why we are using the alternative locution.

<sup>10</sup> For further related discussion of Marr’s theory, including aspects of content that he ascribes to the visual system that could very well be embodied morphologically, see Henderson and Horgan (2000, forthcoming).

<sup>11</sup> Our thanks to participants at the 2009 Bled conference on epistemology for helpful comments and feedback, especially David Henderson, Declan Smythies and Ernest Sosa—and also to Mark Timmons.