





Molyneux's question about perceptual knowledge

Mohan Matthen^a  (mohan.matthen@utoronto.ca)
Jonathan Cohen^b  (cohen@ucsd.edu)

Abstract

Molyneux addressed his question to Locke in two forms. The question that is most often discussed in the literature is the 1693 version—about whether a newly sighted man could distinguish a globe and a cube when they are presented to his sight alone. But in 1688, he asked whether this man could know which was the globe. While Locke and Molyneux probably thought this an unnecessary add-on, we argue that it changes the question. Locke had no account of how one could know a contingent singular fact by sight or by vision. We argue that we know in these ways by using a process of active sensory exploration that employs the modality in question. A new form of Molyneux's question emerges from this understanding of perceptual knowledge. While the variety of sensory exploration methods within and across modalities raises considerable difficulties for a positive answer, the original question, like the later version, can ultimately be answered only by empirical investigation.

Keywords

Common sensibles · John Locke · Molyneux question · Perceptual knowledge · Sensory exploration

This article is part of a special issue on “Molyneux's question today”, edited by Gabriele Ferretti and Brian Glenney.

1 Two Molyneux questions: Perceptual matching and perceptual knowledge

In 1688, William Molyneux wrote to ask Locke whether a newly sighted man visually presented with the two solids could “know which is the Globe and which the Cube? Or Whether he Could know by his Sight, before he stretch'd out his Hand [...]?” The problem Molyneux posed here—call it the 1688 question—was about *perceptual knowledge*: Can this man *know* which solid is which, and know it by sight?

^a Philosophy Department, University of Toronto.

^b Philosophy Department, University of California, San Diego.

Locke did not respond.

Nearly five years later, in 1693, Molyneux, who had by then established a mutually respectful relationship with Locke and was corresponding with him about revisions of his *Essay* for its second edition, tried again. “Whether by his sight, before he touch’d them, he could now distinguish, and tell, which is the Globe, which the Cube.” In this 1693 formulation, the question is once again about *cross-modal identification* or *matching*.¹ But in this second attempt, the 1688 allusion to perceptual knowledge is dropped—here, Molyneux does not use the word ‘know.’

Locke inserted this form of the problem posed by “that very ingenious and studious promoter of real knowledge, the learned and worthy Mr. Molineux”² into the second edition of his *Essay* (§II.9.8).

In this paper, we want to ask about certain issues raised by Molyneux’s first formulation of his question in 1688. It’s natural to think that in the 1693 formulation of the question, Molyneux omits ‘know’ because he thought it raised the stakes to no purpose. His concern was with what were generally thought to be common sensibles—general ideas of shape recognizable by both touch and vision. So, he asks whether the ability to recognize these shapes by touch brings with it the ability to recognize them by vision. The question makes sense only against the background of the assumption that they are common sensibles—if asked with reference to a special tactile quality such as warmth (“Whether by his sight, before he touch’d them, he could now distinguish, and tell, which is the warm, and which the cool”) the question would be absurd. But are these perceptually available shapes *really* common sensibles? Is the idea of a common sensible even tenable? This is one main issue that Molyneux wanted to probe.³

In Locke’s system, general ideas are derived from occurrent perceptual experiences by subtraction. When you have a cube in your hands, you experience many qualities that are irrelevant to its being a cube—its temperature, texture, size, and so on. The general idea, CUBE, is formed by subtracting everything from this experience but the idea of an equal sided solid with six faces. The common sensible

¹ We use the term “cross-modal” in contexts where multiple modalities are in play, but one by one, i.e., not jointly in any perceptual inference or action, and “multimodal” when two or more are used together.

² Locke was using an English spelling of Molyneux’s family name. (For instance, the Wolverhampton Wanderers of the English Premier League play in Molineux Stadium.)

³ Evans (1985) argues that it is a question of a common “concept.” He generally uses this term to denote extra-perceptual abstractions, but since his paper was a draft, we do not know whether he meant it this way. In any event, we claim that the notion that Locke and Molyneux (and Berkeley) are testing is that of a common *sensible*—the product of a generalization specifically from perceptual experience. Evans also claims also that *innateness* is not at issue in Molyneux’s question. In our view, there are so many different questions in play in various treatments of the issue that it is counter-productive to highlight and exclude one issue—especially since Evans is not talking about the historical interpretation of Molyneux’s intention. In particular, the question of knowledge that we take up in this paper is new, and it may actually implicate issues that are not central to other treatments. Nevertheless, Evans’s diagnosis is a powerful antidote against the view that innateness is the main issue involved.

tradition, dominant up to this point, was that these retained qualities were not specific to any one perceptual modality—EQUALITY, for example, is not an idea specific to touch or to vision. Thus, such ideas were held to be apprehensible by vision in the same way as they are by touch. From Locke's point of view, this doesn't make sense. For him and for most other philosophers of his time, all perceptual experience is specific to a single modality. So, if general ideas are created by subtraction from perceptual experience, general ideas too must be modality specific.⁴ There is no such idea as CUBE shared by vision and touch. There is only HAPTIC CUBE, a complex touch-idea generated by subtraction from haptic experience and VISUAL CUBE, which originates from visual experience. (Both these ideas satisfy the geometrical definition of CUBE, but this does not imply that they are the same.) Actual physical cubes instantiate both ideas, but only contingently—the normally sighted person is aware of this by observing and learning the contingent connection. Since the newly sighted man was only able to have a visual experience a few minutes or hours prior, he has had no opportunity to learn this. Thus, Locke's empiricism breaks with the tradition of common sensibles—the beauty of Molyneux's formulation is that it highlights the rupture in sharp relief.

Now if this were all that was at stake in the 1688 question, *knowledge* would seem to be irrelevant. Since knowledge is more stringent than the ability to make a correct identification or match, the mere ability to make a match should be a more significant test: if the newly sighted man cannot even manage that, the alleged common sensibles are in bad shape. So, why not get rid of the mention of knowledge? This is the kind of reasoning that most likely lay behind Molyneux's actual reformulation of the 1688 question and Locke's evident approbation.

As we see it, though, the matter is more ramified than this. For one might ask: How does a blind person come to *know* that something is a cube? Or even that it more closely approximates a cube than a globe? How does a normally sighted person come to know this if the target object is some distance away? Locke and his successors have no good answer to these questions; they have no account of how we might come to know singular contingent facts by perception. In fact, Locke says very little about perceptual knowledge in general—he says only that all knowledge is founded on perception (*Essay* §II.9.15), by which he means only that the *content* of knowledge is built on perception. He does not venture into the question that Plato asked in the *Theaetetus*, which is also the question that the classical Indian

⁴ This inference is dubious. General ideas of shapes are derived from sense impressions by discarding (or disregarding) certain features of the latter. But there is no reason why a sense impression should be wholly modality-specific—a modality-specific impression can, after all, contain some non-specific elements. So, if the general idea of a globe were constructed by discarding all modality-specific features of the *haptic* impressions of these solids, there would be no reason why someone should not be able to construct the same general idea through vision. Of course, this implies that there *are* elements of sense impressions that are not modality-specific—elements like shape and number, for example. Locke and Molyneux dogmatically reject this possibility; but it is unclear how much damage their position would suffer if they were to accept that there is a small non-modal or cross-modal residue that is uncovered in the process of generalization.

schools of epistemology (Matilal, 1986; Phillips & Vaidya, 2024) developed and debated in greater analytic depth than Plato or the European tradition in his wake: How, if at all, can we arrive at knowledge through perception alone?

Knowing a fact is *firmly* and *rationally* believing it—in other words, being confident of its truth, and being so for the right reasons/causes. How can one, starting from a perceptual state, arrive at knowledge, thus defined, about a contingent particular from *perception*? That solid *looks* spherical from here, but is it an ellipsoid? How is one to eliminate this possibility? One could repeat what one did before—look at it again. This would eliminate *some* sources of doubt—for example, “Did I look at the object carefully enough?” But since perceptual impressions are potentially misleading in themselves, some doubt would still remain—I may become more certain that the object looks spherical from here, but it could still be an ellipsoid. The road to knowledge would come to a dead end if nothing is allowed but strings of unconnected momentary perceptual experiences.

In our view, one can best arrive at perceptual knowledge of a contingent singular by an active exploratory process in which the knower interacts purposefully with the target object. Even knowing by touch alone or by vision alone is complex in this way. We will return to elaborate this idea in §IV.

Regardless of what Locke and Molyneux may have thought, this puts the 1688 question in a different light:

Let us Suppose his Sight Restored to Him; Whether he Could, by his Sight, and before he touch them, *know* which is the Globe and which the Cube? Or Whether he Could *know* by his Sight, before he stretch'd out his Hand, whether he Could not Reach them, tho they were Removed 20 or 1000 feet from Him? (emphasis added)

Suppose that there are certain ways to use touch to gain knowledge of shape and also certain ways to use vision for this purpose. Before he was able to see, the newly sighted man was able to use these methods to *know* whether an object in his hands is a sphere or a cube. The questions suggested by this reading of the 1688 question go along the following lines: Is there a sensible quality, vision-based or not, about which touch-based ways of knowing would guide him in using *vision* to know? For example, Berkeley concedes that “the visible square is more fit than the visible circle to represent the tangible square [...] because the visible square does, whereas the visible circle doesn’t, contain several distinct parts by which to mark the several distinct corresponding parts of a tangible square” (NTV §137). Should Berkeley allow, then, that these “corresponding parts” could be a guide to vision-based ways of knowing that something was a square, even for the newly sighted man? Of course, Berkeley puts the argument forward in connection with the 1693 question. But it has equal application to the question of ways of knowing. So, the question is this: Would he be able to use vision-related methods directly and immediately—innately or by some kind of transfer from his haptic abilities? Or must he acquire the vision-specific ways of perceptual knowledge by trial and

error? This is a different question than the one about the modality-specificity of the retained ideas in CUBE and SPHERE.

To state our position more explicitly: the 1693 question is about the identification of shapes. The newly sighted person can identify shapes by touch. Is she also able to identify shapes by sight? If she can, then either she has innate knowledge of a common sensible, or this knowledge can be transferred from her haptic ability. This is the question that has been investigated in most of the literature so far. The 1688 question takes us in another direction. There are certain methods that perceivers use when they seek to *know by touch* what shape something is. There are also methods that normally sighted perceivers use when they seek to know *by vision* or *by touch and vision* what shape something is. Now, consider a blind person who has access to the first set of methods. Suppose that vision is surgically or otherwise bestowed upon her. Does she have access to the later mentioned methods? Is she able to know *by vision* what shape something is? Assuming that *knowing* that something is a cube doesn't simply recapitulate the perceptual identification of this thing as a cube, then the questions are different. Our aim is to demonstrate that there are ways of knowing by perception that are not simply a reduction of uncertainty by touching or looking at the object again.

In the next two sections, we will prepare for the consideration of perceptual knowledge first by identifying and discussing a secondary question that is irrelevant to us (§§II-III). We will then introduce the idea that perceptual knowledge about the external world can be arrived at by multimodal perceptual exploration (§IV). With respect to this route to perceptual knowledge, and pushing the secondary question to one side, we will argue that there is no *a priori* answer to the 1688 question (§§IV-VI). This reinforces our earlier published contention (Cohen & Matthen, 2021; Matthen & Cohen, 2019) that there are many Molyneux questions, not just one, and that most of these must be tackled empirically, not *a priori* as many philosophers, starting with Locke, approach the Molyneux problem.

2 Developmental delays in cross-modal perceptual matching

We will now explain the secondary question that complicates the proper understanding of both Molyneux questions by reference to a famous recent research program that, like many recent philosophical and psychological discussions, focused on Molyneux's 1693 question. In particular, this later question (rather than the question of knowing) is what Pawan Sinha and colleagues investigate in their celebrated study (Held, R. et al., 2011). These authors begin by noting that, as they put it: "A few studies of cross-modal matching by neonates have reported that they are able to visually choose between two objects that they have previously felt only via touch." They dismiss these results, however, as "hard to replicate" and offer evidence that they take to be suggestive of a negative answer to Molyneux's question.

Five patients who had been profoundly blind since birth with congenital cataracts gained sight by cataract surgery when they were adults. “As soon as was practical after surgery,” these patients were tested on a series of pairs of objects, each haptically presented and then presented again either haptically or visually. They were asked whether these were the same or different. These patients were reliably successful when the objects were both presented haptically, and after two days, when both objects were presented visually, but for a few more days, successful only at chance levels at the cross-modal task. After about a week, however, these patients were fully up to the performance levels of people who had been sighted since birth. In view of these findings, Held et al write: “Our results suggest that the answer to Molyneux’s question is likely negative.” They deem the patients’ failures of the first few days to justify a negative answer to the Question as it was put in 1693. However, they realized that since these patients spontaneously acquired the ability to make these matches in a very short time, this answer must be significantly qualified.

It will be useful for what follows to distinguish two kinds of process that might, in real life, so to speak—as opposed to the idealized worlds created by philosophical thought experiments—cause a delay between the restoration of the eyes and the acquisition of the ability to discern and identify (or know) shapes well enough to match them to shapes earlier known by touch.

One possible cause of such a delay is time taken for the patient’s eyes and visual system to come to full working order, and for the patient to become accustomed to using them optimally. Immediately after sight has been restored, patients might have trouble focusing on a particular object, or scanning a scene to find it. Their visual systems may have trouble segmenting the scene, they may have trouble focusing attention on the right things, and so on. In short, the newly sighted person might need some experientially based trials to get beyond the “booming, buzzing confusion” of newfound vision and get properly going with the proper use of this sensory faculty. Call this a *developmental* delay. This kind of delay has nothing to do with comparisons with touch; it is just about getting accustomed to vision itself.⁵

A more germane cause of delay is *cross-modal mismatch (cmm)*. One kind of *cmm* occurs when the existing modality provides content that the restored modality does not: for example, touch provides haptic texture, which relates to high frequency ups and downs on the surface—the kind of roughness or smoothness that one can feel—while vision gives us information about visual texture—high frequency variance in reflectance on a surface. While each of these properties of a surface predicts the other quite well, and each is therefore used as a proxy

⁵ This, presumably, is the point Evans (1985) was making when he said, “Molyneux’s Question is about whether a born-blind man who can see a circle and a square would extend his concepts to them. It is not a question about how soon after the operation, and by what process, a newly sighted man would be able to see.” (ibid. 365–366). (The phrasing is unfortunate—sighted people are able to see—but the intended point is correct.)

for the other, the correlation is contingent. In other words, haptic roughness and smoothness is not the same property as visual roughness and smoothness. Call this *content* mismatch. Another kind of *cmm*—call it *content-presentation* mismatch—occurs when the impressions provided by the existing modality are dissimilar in psychologically important respects (form, encoding, access, etc.) to those provided by the restored modality even when their content is the same. For example, suppose a first modality represents locations of distal items by contact and a second, newly acquired, modality represents them distally; then, even if the two modalities present the same content, it may be that experience is required to establish the match in content.⁶

Clearly, Molyneux's intention was to raise the question of cross-modal mismatches between vision and touch. His insight was that as far as Locke's system is concerned, intermodal content presentation mismatches are *content* mismatches. For, as recounted earlier, the phenomenal difference between touch and vision implies that the general ideas derived from touch and vision are different—neither gives us the general idea CUBE; each gives us a modality-specific idea. From this point of view, nothing more needs to be said about knowing by perception that something is a cube—if the newly sighted man lacks the idea of a visual cube, then there is nothing he can do to know by sight that something is a cube. And this has nothing to do with developmental shortcomings. We can put these aside.

Given that Molyneux meant to be asking about a delay due to a cross-modal mismatch—specifically, content-presentation mismatch—it is inappropriate to answer his question in the negative based on a developmental delay: for example, an inability to focus the eyes due to a lack of experience (or, for that matter, a difficulty experienced in opening them). But this is how a number of philosophers in fact reason (mainly French *philosophes*, according to Degenaar et al. (2024, sec. 3)). Thus, consider Merleau-Ponty's position as reported by Shaun Gallagher:

In the initial visual perception for the Molyneux patient [...] 'everything is at first confused and apparently in motion. Discrimination between coloured surfaces and the correct apprehension of movement do not come until later, when the subject has learned "what it is to see" [...] It thus appears that the empirical cases of congenitally blind subjects who gain vision, exceptional though they are, help to show the importance and necessity of [repeated] experience for the perceptual process. (Gallagher, 2005, p. 156)

Merleau-Ponty (who Gallagher seems to endorse) is apparently saying that Molyneux's question should be answered negatively because developmental delays render the newly sighted patient incapable of visually discerning shapes—specifically, it takes a while before the subject can learn how (or rather “what it

⁶ Similar worries concerning cross-modal differences in spatial representation will arise if Batty (2010) there are important differences in the spatiotemporal structure of perceptual experience connected with different modalities.

is”) to see. But this, we would argue, overlooks what is at issue in Molyneux’s formulation.⁷ Molyneux did not ask whether the visual system might take time to function properly after the opaque lenses of the eyes were removed, thus permitting light to be fall on the retina. He was asking about the newly *sighted* man, not the recently surgically operated one. His question was whether a person would be able to know or identify shapes by *sight* that were previously familiar through touch when these shapes were visually presented for the first time. The question assumes that *sight*, not just the function of the eyes, has newly been restored, and it is thus completely independent of any imperfection of the visual system or of the way it is used immediately after restoration of the eyes.

It is worth noting that Sinha and colleagues’ tentative (and widely reported) negative verdict with regard to Molyneux also seems to rest on developmental delays. His patients were never taught, and did not have the opportunity to learn, the correlations between the familiar haptic presentation of various solids and their visual counterparts. So, it seems that their ability to match these presentations emerged spontaneously without systematic teaching or learning around day 5 after surgery. The predictable time-course of this emergence despite assumed differences in individual developmental histories suggest operational delays—that is, they suggest that the restoration of sight lags the removal of congenital cataracts by a few days. But if the initial delay was not traceable to a cross-modal mismatch, then these findings do not support a negative response in the sense intended. Indeed, as Sinha’s cohort is fully aware, they serve as the gateway to further investigations that establish a positive answer. At the very least, they demonstrate that the performance of patients immediately after cataract surgery is not a good test of Molyneux’s question because developmental delays in the recovery of vision falsely intimate a content-presentation mismatch.

A similar point can be made about perceptual knowledge. We have certain ways of knowing the shape of an object we perceive. Some of these ways are specifically haptic, some specifically visual, some multimodal. Normally sighted adults have access to these methods—let’s say that they are fully possessed of *perceptual-epistemic* abilities. We can ask: does the newly sighted person have access to the visual and the vision-involving multimodal epistemic abilities? In approaching this question, one must distinguish between developmental delays and delays due to cross-modal mismatch. It could be that there is a psychologically relevant difference between vision and touch that makes the touch related methods of knowing inapplicable to vision. And the newly sighted person may for this reason have to learn how to use vision to acquire knowledge. Or it could be that the newly sighted

⁷ Gallagher (*ibid*) clearly distinguishes the processes we identify above: “At first, perception is confused,” he says, but empiricists maintain that even after this confusion is overcome, the sense modalities do not “communicate” or “educate” each other. However, we diverge from him in treating the initial confusion of perception to be irrelevant to the intent of Molyneux’s question as posed to Locke.

person is already capable of using vision to acquire knowledge, though she needs some experience before this ability is fully realized.

Molyneux's 1688 question is about perceptual-epistemic abilities involving the absent modality due to cross-modal mismatches. It is not about developmental delays. Any empirical investigation of epistemic abilities should have some way of making this distinction.

3 Retinal information and visual experience

There is a further complication about multimodal perceptual experience. (This complication mostly eluded Locke, as well as many later writers; Berkeley notices it, but sidesteps, most likely because it threatened the empiricist foundations of his philosophy.) The problem is this: At what level of combination can qualities presented by one modality be mimicked by another? Let us explain.

Under the influence of the geometrical optics tradition stretching from ibn-al-Haytham to Johannes Kepler, Locke apparently assumed, and Descartes explicitly held, that visual experience corresponds point by point with the retinal projection from the scene it represents. Empiricist philosophers seem to generalize this assumption, though not explicitly—perceptual experience in each modality structurally matches the proximal data in that modality's receptor array with no grouping, segmentation, filtering, or other embellishments. If this were true, Molyneux's question would trivially yield a negative answer because no two modalities share the same proximal data.⁸ This is one premise that lies behind Locke's confident approach to the question. But the assumption can, of course be questioned—and, in modern perceptual science, it generally is.

We can illustrate this by focusing on the case of vision, i.e., on the substantial gaps between the retinal projection and visual experience. To begin, even putting aside the cortical processing that the retinal image undergoes, the spatial organization of the retina differs substantially from that of visual experience. The spatial resolution of the retina is non-uniform (it has a higher concentration of cone cells at the center than the periphery); even at the center, its resolution is limited by the packing density of photoreceptors; in the periphery, the retinal projection is colour-desaturated and lacking in detail—its main function appears to be the detection of change and movement. The retina has a blind spot where it is attached to the optic nerve; blood vessels cast shadows on the retinal projection; and the projection is constantly shifting abruptly due to constant saccadic motion.⁹ None

⁸ This is a restatement of the proposition we asserted in Matthen & Cohen (2019): "Molyneux's question is inapplicable to point-data: colour (/intensity) is the only visual point-datum, and [...] it cannot be cross-identified with any tactual point-datum. However, the question is compelling when applied to the transfer of integrated constructs from one modality to another." (*ibid.* p. 48)

⁹ For discussion of these and other disanalogies between the retinal projection and the visual field, see Hardin (1988, Chapter 1). Of course, the retinal projection idea was already opposed, quite vehemently, by the Gestalt psychologists, and following them, by Merleau-Ponty.

of these features shows up, as such, in normal visual experience (nor, therefore, in normal visual knowledge) of tables and chairs and the like. On the contrary, our visual experience seems spatially uniform, potentially unlimited in resolution, unblurred throughout its extent, without holes or endogenously generated shadows, and stable over time.

For that matter, and as the entire post-Helmholtzian tradition in computational vision attests, the visual system builds and feeds to visual experience a series of representations that move well beyond the information available at the retina. As this tradition has emphasized, the visual system deploys a suite of computational techniques for recovering the stable object properties presented in visual experience (e.g., color and form) from retinal data that conflate these properties with contributions from the perceptual situation (e.g., illumination and viewing angle). In order for visual experience to present such stable object properties as color and form (as it obviously does), it cannot merely mirror the receptor data; rather, the visual system must subject those data to a series of non-trivial computations whose outputs to experience and knowledge differ significantly from their receptor inputs (Cohen, 2010, 2012).

This matters to our question because, as noted, every one of these respects in which visual experience comes apart from the proximal data present on the visual receptor array amounts to a potential respect in which enabling the operation of the receptors in a hitherto blind person is, by itself, insufficient for restoring visual experience and visual knowledge. As we emphasized in the previous section, surgically enabling the eyes is different from making a patient see.

Interestingly, Berkeley was aware—in a way that Locke seems not to have been—that the structure and features of visual experience cannot be assumed to be delivered trivially from the data delivered to the retinas/eyes, but that, instead, these are constructed from those data by non-trivial processing. In his *New Theory of Vision*, published in 1709, twenty-one years after Molyneux's first question, he writes:

Whenever we look carefully and in detail at an object, successively directing the optic axis to each point on it, the motion of the head or eye traces out certain lines and shapes that are really perceived by feeling but so mix themselves with the ideas of sight (so to speak) that we can hardly avoid thinking of them as visual. (§145)

The claim, in short, is that sighted people become aware of certain spatial properties (such as lines too long to take in with a single glance) by virtue of “the motion of the head or eye,” and that these properties are, consequently, “perceived by feeling” (or in other words, by touch, in which he includes proprioception). So, vision and “feeling” are phenomenally intertwined—we can “hardly avoid” thinking of properties “really perceived by feeling” as visual.

One would expect Berkeley, extreme empiricist that he is, to attribute this to learning, since he follows the tradition in taking the modalities to be completely

insulated from one another. Indeed, he does say that the blending of visual and tactual ideas is by a “habitual connection” similar to that which we observe with semantic convention (§147). And accordingly, his negative answer to the Molyneux question of 1693 is emphatic and unequivocal (see §§132–135; also §110). However, he does clearly see, and appears to be tempted by, some of the attractions of the alternative. For he says that our ideas are a “universal language of the Author of nature” designed to highlight associations “so as to get the things we need for the preservation and well-being of our bodies and avoid whatever may be hurtful and destructive of them.”¹⁰ Thus, it is no accident that “the motion of the head or eye” serves as a proxy for tactual ideas of distance and direction, which in turn get pasted on to visual ideas.

The wonderful art and contrivance with which it is fitted to the goals and purposes for which it was apparently designed; and the vast extent, number, and variety of objects that are at once suggested by it with so much ease, speed and pleasure; these provide materials for much speculation—pleasing speculation—and may give us some glimmering, analogous prenotation of things that we can't properly discover and comprehend in our present state. (§148)

To identify or know a sphere, especially one that is sufficiently large/close to occupy a significant part of the visual field, the eye must saccade and change focus, and the perceiver must actively scan its surface. More generally, it must construct a scene by gathering information; it must also be an instrument that can be directed by the perceiver to answer her queries.

Now, consider this intertwining of vision and touch in the context of our earlier remarks about the distinction between developmental delays and delays due to cross-modal mismatches in the visual performance of the newly sighted, whether these newly sighted individuals are newly born or newly operated upon. After the eyes/visual transducers are established in working order, it takes time for sight to be established. The question is whether (or how much of) this delay is due to *associationist learning*, as empiricists like Berkeley would have it, as contrasted with *developmental acquisition*, a natural process triggered by experience, but not an increasing function of past experience (as associative learning is). We have argued that the time-course of developmental processes that begin with the data on visual receptors is one cause of a delay between the enabling of the eyes and the acquisition of vision. Merleau-Ponty insightfully points this out, but assumes, wrongly in our view, that it is directly relevant to Molyneux's questions. The reason he is mistaken is that these delays are plausibly thought of as a developmental, rather than content-mismatch, issue. The matter can be empirically adjudicated: the emergence of the above-stated inter-modal skills on a fixed timetable (akin to the one-week delay in the patients who were operated on by Pawan Sinha's group) would

¹⁰ For discussion of perceptual learning in Berkeley and its relation to linguistic meaning, see Copenhaver (2014).

be evidence of developmental acquisition; sensitivity to quantity of data to which a patient is exposed would be evidence of learning. Thus, as we have contended in earlier work (Cohen & Matthen, 2021; Matthen & Cohen, 2019), Molyneux's Question is not answerable *a priori*.

4 Perceptual knowledge and sensory exploration

Return now to Molyneux's 1688 question concerning (not just identification or recognition, but) whether the congenitally blind subject with sight newly acquired "Could, by his Sight, and before he touch them, *know* which is the Globe and which the Cube?" What does it take to *know by visual perception* (and not, for example, by laser-aided measurement) which was the cube and which the globe? An important part of the answer, we'd like to suggest, centers around the notion of *sensory exploration*.¹¹

To illustrate, we begin with a unimodal example. Consider a normally sighted person who visually apprehends a three-dimensional object in a quick initial glance. Its visual appearance gives her a *prima facie* reason to believe that it is spherical (Burge, 2003; cf. Pryor, 2000). Of course, neither this visual appearance nor the ensuing visually based belief is conclusive: the spherical look of the object may be a trick of the viewing angle, shape contrast, etc. As well, it is incomplete: though something may look like a sphere when viewed from one angle, it arguably doesn't really look spherical until one has viewed it "in the round." (This is crucial for irregular solids.) For these reasons, if the stakes are sufficiently high, our sighted person may wish to examine matters further, engaging in a process of deliberate investigation or exploration (Matthen, 2012). For example, as a way both of gaining a more comprehensive view of the object in three dimensions and of allaying the sources of doubt mentioned, she may wish to examine the object's visual appearance in a variety of viewing angles and distances. The visual appearances gleaned in this wider range of conditions enables her to form a more complete impression; moreover, they can supply further evidence that helps (when combined with suitable background beliefs about the stability of the perceptual conditions, etc.) allay sources of doubt regarding the initial belief concerning the object's form. They do this both by supplying quantitatively additional data that complete and corroborate the initial appearance—looking at something twice or for a longer time is better than looking at it once or for a shorter time—and also by supplying qualitatively different data — e.g., appearances from distinct viewing angles — that, taken collectively, control for possible confounds not eliminated by the initial appearance taken by itself.

Sensory exploration can also be multimodal. If our normally sighted person picks up the object and turns it over in her hands, she may arrive at a highly corroborated belief concerning shape based on both touch and vision in a mutually

¹¹ For discussion of sensory exploration and its epistemic properties, see Matthen (2012).

reinforcing way. Again, the new data obtained by haptic exploration can support the initial visual appearance both quantitatively—they add to the stock of evidence supporting conclusions about the object's form—and qualitatively—the newly obtained haptic data are not vulnerable to confounds of viewing angle (though they may be vulnerable to different threats; see below). Note that, as observed in §III above, such reinforcement can apply only to processed properties. Point data presented by touch are never comparable with point data presented by vision. One can tell that something that looks spherical, for example, also feels spherical. But touch-based point data (something on the intense/faint pressure spectrum) cannot confirm vision-based point data (colour).

Such unimodal and multimodal explorations can, if they are deliberate, attentive, and self-aware, eliminate doubt-producing alternatives and build toward the well-founded epistemic confidence required for knowledge. They provide mutually reinforcing evidence that makes it increasingly probable that the resultant visual belief was accurate, well-founded, and not in need of further adjustment or revision.¹² In the unimodal visual case, for example, the yellow visual appearance of a banana seen from a different angle and distance provides corroborating evidence regarding its color (making it improbable that its initial yellow appearance was merely a trick of the illumination). In the multimodal exploration of shape, the spherical visual and haptic appearance of a billiard ball provides corroborating evidence regarding its form (making it improbable that the initial spherical visual appearance was merely a result of noise, confusion with other objects, or simply insufficient experiential data); and so on. Using these methods, the perceiver approaches empirical certainty about the object's color and form, and, thus, ultimately, that the object really is a yellow banana/spherical billiard ball, not some simulacrum thereof.¹³

The observation that sensory exploration can, ubiquitously and in quite ordinary cases, provide evidence that shores up the warrant of perceptual beliefs generalizes to a wide range of perceptible qualities in a wide range of perceptual modalities. After all, it is a permanent feature of our epistemic lives that the circumstances of perception are incomplete and noisy, that the energy striking our transducers is the joint product of perceived objects and sundry parameters of the perceptual conditions (and of our own bodies), and that perceptual signals there-

¹² Some writers (e.g., Cohen, 1999; DeRose, 1992, 2009) have held that thinkers can qualify as having knowledge in low-stakes scenarios even if they have done less to meet skeptical threats than would be required for them to have knowledge were the stakes higher. Extending this view in the natural way to cases of perceptual knowledge, one might hold that thinkers in low-stakes scenarios who have not, or not yet, engaged in the sensory exploratory activity needed to allay possible sources of doubt nonetheless count as having perceptual knowledge. We need not take a stand on this question; what we say here will still have force so long as it is true, as we contend, that there are ordinary cases in which perceptual knowledge demands sensory exploration.

¹³ Matthen (2012) argues that these methods lead to “empirical certainty,” i.e., certainty about the facts specific to the case, as opposed to certainty regarding the subject's epistemic detachment from the empirical world and all of its situations and facts.

fore underdetermine just what condition of the distal world is responsible for the perceptual experience. Sensory exploration allows us to overcome these obstacles to perceptual knowledge by affording a wider class of perceptual evidence that eliminates various potential defeaters of our perceptually informed beliefs — defeaters not excluded by our initial perceptual encounter, taken all by itself.

5 The specificity of sensory exploration

Knowledge by sensory exploration provides a way into Molyneux's relatively underdiscussed 1688 question about the newly sighted man's capacity to know shape "by his sight." For now we can ask whether and in what ways a Molyneux patient expert in the use of sensory exploration to achieve haptic knowledge will be similarly adept in the use of sensory exploration for *visual* knowledge (once she begins properly to see). Note that this is a question about unisensory knowledge—does the newly sighted man know how to know *by his sight*. We want to offer a general reason for skepticism about a direct and complete transfer of these epistemic capacities from one modality to another: for the forms of sensory exploration that contribute to perceptual knowledge are plausibly local to particular sensible qualities in particular sensory modalities.¹⁴ But it should be noted that our sceptical answer does not cover multimodal exploration, which uses data from visual exploration to *corroborate* haptic data without specifically ruling out haptic defeaters. Molyneux did not ask about multimodal exploration, and we follow suit.

Let us begin by looking at cases where the acquisition of knowledge by sensory exploration is *intra*-modal. For example, sensory exploration supporting the visual apprehension of shape typically involves, among other things, manipulating a seen object or changing one's position with respect to it so as to view it from different angles. Crucially, this form of sensory exploration aids the visual apprehension of shape because of how the visual presentation of shape *in particular* varies with changes of viewing angle. This presentation depends (as Locke and Molyneux knew from Kepler's ophthalmological optics) on the two-dimensional projection of three-dimensional shapes on the two-dimensional retina. Different solids project identically to the retina from the same angle, and the same solid projects differently from different angles. Thus, the epistemic recovery of three-dimensional shape by visual exploration is governed by projective geometry.

Obviously, however, projective geometry does not govern the presentation of other qualities we apprehend by vision, and so visual exploration of how these other qualities vary is subject to different rules and methods. For example, in the

¹⁴ Of course, it may be, that some forms of exploration also support the matching task of Molyneux's 1693 formulation; thus, Schwenkler (2012) suggests that the matching task is best tested by allowing the newly sighted subject to move around the object, and so to assemble depth cues not available in instantaneous visual presentations (but cf. Connolly, 2013). (Neither of these authors distinguishes the matching question from that concerning perceptual knowledge that we address here.)

context of uncertainty about the color of an item, exploration involving varying the visual angle may not be helpful (though in some cases it may be). Rather, in this case, it is generally useful to gain visual experience of the color of a known distinct item under the same illumination (Hurlbert & Ling, 2005) or of the same item under a known distinct illuminant, thereby visually gathering information that allows one to disentangle the contributions to the undifferentiated proximal signal made by the object and the illumination, and so to decide between distal alternatives that present the same signal to visual transducers at a moment. (Is it, say, a white sphere under red illumination or a red sphere under white illumination?)

The lesson from this pair of examples is that, though they both involve visual perception, and though it is equally true of them that sensory exploration enhances our epistemic positions by ruling out possible defeaters, it does so in different ways. This is because, as it happens, distinct defeaters threaten visual perceptual knowledge of shape and color, hence distinct classes of additional perceptual evidence are useful in ruling out these threats.

Now, the same is true of the intermodal context raised by Molyneux with respect to even the same property, shape, and this suggests a complication of our discussion in §IV. For sensory exploration undertaken in a second modality will not, in general, rule out the potential defeaters specific to a first when the two are directed on a common sensible. We have said that if something initially looks like a sphere, a perceiver can improve her epistemic position with respect to its shape by changing her angle of view. But this is not the form of sensory exploration that supports her haptic knowledge of shape. Haptic exploration for shape would include the quite distinct activity of rotating the object in her hands to test for symmetry in all directions. That there is this difference in the forms of exploration relevant to apprehension of form in vision and touch is unsurprising; it results from the fact that the presentation-variability of shape in vision is different from that in touch, and thus that the potential defeaters threatening knowledge of shape is distinct in the two modalities. So, the types of exploration needed to overcome visual defeaters would be expected to differ from those required for haptic defeaters.

Now, though haptic exploration does not provide evidence that speaks to the potential defeaters of visual experience/belief, it is not for that reason epistemically worthless. After all, when conditions are propitious, haptic exploration can, by providing evidence undercutting potential defeaters of *haptic* experience, secure *haptic* perceptual knowledge of the item's form. And, indeed, if this information corroborates one rather than other of the distinct distal alternatives left open by your initial visual experience, you might even count following haptic exploration as possessing *knowledge* (even *perceptual knowledge*) of its form. But it seems right to think of the case as one in which haptic perceptual knowledge aligns with a distinct and epistemically weaker state informed by visual experience, rather than

one in which haptic exploration raises the outputs of visual perception to the level of *visual knowledge* in particular.¹⁵

Of course, in saying that we cannot assume that sensory exploration in a second modality *will* rule out the potential defeaters specific to a first, we do not mean to encourage the assumption that sensory exploration in a second modality *can't* do this. Everything depends on the particulars of the modalities, the potential defeaters specific to them, and the exploratory procedures available to the perceiver. This is not a matter to be answered from the armchair, but through case by case empirical examination.

This conclusion motivates a somewhat nuanced response to Molyneux's 1688 question about the transfer of perceptual knowledge from one modality to another. If we are right in the foregoing, the achievement of perceptual knowledge of form by vision (at least in sufficiently high-stakes scenarios) requires the use of sensory exploration to rule out potential defeaters. Moreover, the specific epistemic gains provided by visual sensory exploration —the gains required for the status of knowledge of form by visual perception— are, in general, not provided by sensory exploration in other modalities such as touch, even if those other modalities share form as a common sensible. But if so, then the tactile expertise of the man newly restored to vision is unlikely to help him achieve *visual* perceptual knowledge of form. This man's tactile expertise may include expertise with *tactile* exploration of form; moreover, he may be able to use that expertise to confirm initial visual appearances of form, and perhaps even achieve on this basis knowledge of form — even *perceptual knowledge* of form. However, insofar as the man's tactile expertise fails to deliver the epistemic goods specific to the *visual* exploration of form — i.e., the power to eliminate the specific defeaters threatening the epistemic power of the visual apprehension of form, in particular, he would not through his tactile expertise be in a position to do what is needed to convert his initial visual experience into *visual knowledge*.

This leaves two possibilities. The first is that the Molyneux patient is simply unable to achieve visual knowledge until he learns to do so by some painstaking process of experiential conditioning. The second is that some set of knowledge-gaining procedures is available to her (either as a result of her innate endowment or developmental learning) facilitate this visual knowledge without experiential conditioning.¹⁶ This may seem implausible at first sight. But reflect on a conception of vision not as a purely receptive faculty, but rather one that includes the exercise of interactive exploration. Perhaps if vision is active in this way, sensory exploration would be included in some "ways of seeing." (This would provide grounds

¹⁵ Compare: you might gain further or more conclusive evidence of the item's shape from a non-perceptual source such as memory or testimony; your knowledge of its shape would then be memorial or testimonial rather than visual, even if it aligned with one of the alternatives initially made available by visual experience.

¹⁶ As we remarked with respect to a different form of developmental delay in §II, if it turned out that these capacities were not immediately manifest, we would not take this point as warranting an interesting negative answer to Molyneux.

for a positive answer to Molyneux's 1688 query "Whether he Could know by his Sight".) Again, it is impossible to decide between these alternatives on an *a priori* basis: further empirical evidence is needed.

6 Conceptual knowledge of unfamiliar modalities

Before concluding, we want to consider the possibility of an alternative, non-perceptual, route to visual knowledge suggested by the case of Esref Armagan, a congenitally blind man who is a gifted painter. Armagan has been profoundly blind since birth. His deficit is not merely a cataract that blocks the light entering an otherwise functional eye—in his case, "one eye is absent and the other is a microball that has no light sensitivity" (Kennedy & Juricevic, 2006a, p. 507). Despite these visual deficits, Armagan draws and paints landscapes with realistic visual detail, including colour and distance cues such as foreshortening and perspective. According to a press report:

His father, Nazim Armagan, began to introduce him to various objects and teach him concepts like roundness, sharpness, etc. "To get to know shapes and objects, a blind person must be able to hold it in their hands," Armagan said, explaining that it was the only way to perceive the object from all six directions – top, bottom, and so on. "If the object cannot be wholly examined at once, then the picture in the mind's eye is disconnected, incomplete and inaccurate," Armagan added. So, when it came to things he could not hold, his father would give Armagan models. It all began with a butterfly, when his father gave Armagan a paper butterfly to understand what the insects were shaped like. While holding it, Armagan had the idea that he could try to draw the shape – a test to see whether he accurately perceived how objects looked. He placed the paper on a surface that caved in under his pencil, using the relief method, so he could perceive the drawing through his fingertips, just like the model, and compare them with one another. Later, when he began painting on canvas, Armagan would use a sticky rope to create the outline of his paintings so that he could feel the lines. "I have to feel what I am drawing with my fingertips because that's how I see," Armagan explained. (Balkiz, 2022, press-style paragraphing eliminated.)

Note that despite Armagan's claim that haptic exploration is "how I see," foreshortening is not given by touch: rather, he was *taught* that things that were further from him had to be drawn smaller to look the same size.

Kennedy and Juricevic, who have studied such abilities in more than one blind individual, found that Armagan's skills were remarkably developed (cf. Kennedy, 2003; Kennedy & Juricevic, 2003, 2006a). They asked him to draw "solid and wire

cubes, in several positions—a cube directly in front of him, a cube moved to the left, a cube moved to the left and down,” and also a cube balanced on a point with a vertex pointing toward him. The resulting drawings shows convergence and foreshortening appropriate to one-point perspective, with the occluded lines eliminated appropriately in the drawings of the solid cubes. Kennedy and Juricevic conclude that certain descriptive concepts have visuo-haptic transfer:

The blind and the sighted often hear that things in the distance are pictured small, and look small. Esref reports being told that roads converge in pictures. Esref’s drawings may be applying this principle to cubes, combining it with foreshortening. (Kennedy & Juricevic, 2006b)

Obviously, Armagan has no *visual* knowledge: he doesn’t have vision. However, he has evidently acquired *conceptual* knowledge of how three-dimensional shapes are visually presented—knowledge by description, to use Russell’s (1911) phrase. The considerable extent of this knowledge is instructive, for it gives us a clue of how a Molyneux subject may have been able to acquire the knowledge needed visually to distinguish shapes at first presentation. It might very well be that if Armagan were miraculously to acquire vision, he would (after the requisite operational delay) be able to distinguish a cube from a sphere because he would be able to count the sides and edges of the cube and discern the effects of foreshortening on each. This would confirm and supplement Jonathan Bennett’s (1965) conjecture that a Molyneux subject would recognize the cube by the fact that it has vertices, faces, equal sides etc. while the globe does not. Armagan’s knowledge of the visual cuts against the idea that content-presentation mismatches can only be overcome by direct experience, and so by sensory exploration building from that experience.

To summarize: Armagan’s abilities suggest that it is, at least, possible that a man a blind man could know what a cube and sphere look like and that he could use this knowledge to discern the two shapes by sight if he were to acquire the power of vision. However, they do not demonstrate that a newly sighted person could *know* by sight which was which.

7 Tackling the 1688 question

We have distinguished two questions regarding perceptual knowledge of shape as a common sensible. The first regards knowing by sight. We argued that since methods of knowing for different sensory properties were specialized, it was unlikely that methods of knowing tactual properties would simply transfer over to vision. Given this, a ‘yes’ answer to Molyneux’s 1688 question would have to rest on the availability of innate or acquired methods of sensory exploration yielding knowledge *by sight* of shape. The second question regards knowing by perception. Could we have innate knowledge of visual exploration procedures that would reinforce or bolster knowledge by touch? Here, a ‘yes’ answer is less difficult to conceive of *a priori*, but again a final answer can only be obtained by empirical investigation.

Acknowledgments

We are grateful to Matthew Fulkerson, Huadian Liu, and Andy Sin for helpful discussions that have improved this paper considerably.

References

- Balkiz, K. N. (2022). Esref Armagan: A blind Turkish painter who sees through his fingertips. *TRTWorld*.
- Batty, C. (2010). Olfactory experience I: The content of olfactory experience. *Philosophy Compass*, 5(12), 1137–1146. <https://doi.org/10.1111/j.1747-9991.2010.00355.x>
- Bennett, J. (1965). Substance, reality, and primary qualities. *American Philosophical Quarterly*, 2(1), 1–17.
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, 67(3), 503–548. <https://doi.org/10.1111/j.1933-1592.2003.tb00307.x>
- Cohen, J. (2010). Perception and computation. *Philosophical Issues*, 20(1), 96–124. <https://doi.org/10.1111/j.1533-6077.2010.00185.x>
- Cohen, J. (2012). Computation and the ambiguity of perception. In G. Hatfield & S. Allred (Eds.), *Visual experience: Sensation, cognition, and constancy* (pp. 160–176). Oxford University Press.
- Cohen, J., & Matthen, M. (2021). What was Molyneux's question a question about? In G. Ferretti & B. Glenney (Eds.), *Molyneux's question and the history of philosophy* (pp. 325–344). Routledge.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89. <https://doi.org/10.1111/0029-4624.33.s13.3>
- Connolly, K. (2013). How to test Molyneux's question empirically. *I-Perception*, 4, 508–510. <https://doi.org/10.1068/i0623jc>
- Copenhaver, R. (2014). Berkeley on the language of nature and the objects of vision. *Res Philosophica*, 91(1), 29–46. <https://doi.org/10.11612/resphil.2014.91.1.2>
- Degenaar, M., Lokhorst, G.-J., Glenney, B., & Ferretti, G. (2024). Molyneux's problem. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Summer 2024). <https://plato.stanford.edu/archives/sum2024/entries/molyneux-problem/>; Metaphysics Research Lab, Stanford University.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929. <https://doi.org/10.2307/2107917>
- DeRose, K. (2009). *The case for contextualism: Knowledge, skepticism, and context, vol. 1*. Oxford University Press.
- Evans, G. (1985). Molyneux's question. In G. Evans (Ed.), *Collected papers*. Oxford University Press.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford University Press UK.
- Hardin, C. L. (1988). *Color for philosophers: Unweaving the rainbow*. Hackett.
- Held, R., Ostrovsky, Y., de Gelder, B., Gandhi, T., Ganesh, S., Mathur, U., & Sinha, P. (2011). The newly sighted fail to match seen with felt. *Nature Neuroscience*, 14(5), 551–553. <https://doi.org/10.1038/nn.2795>
- Hurlbert, A., & Ling, Y. (2005). If it's a banana, it must be yellow: The role of memory colors in color constancy. *Journal of Vision*, 5, 787–787. <https://doi.org/10.1167/5.8.787>
- Kennedy, J. M. (2003). Drawings from Gaia, a blind girl. *Perception*, 32(3), 321–340. <https://doi.org/10.1068/p3436>
- Kennedy, J. M., & Juricevic, I. (2003). Haptics and projection: Drawings by Tracy, a blind adult. *Perception*, 32(9), 1059–1071. <https://doi.org/10.1068/p3425>
- Kennedy, J. M., & Juricevic, I. (2006a). Blind man draws using diminution in three dimensions. *Psychonomic Bulletin & Review*, 13(3), 506–509.
- Kennedy, J. M., & Juricevic, I. (2006b). Foreshortening, convergence and drawings from a blind adult. *Perception*, 35, 847–851.
- Martin, M. (1992). Sight and touch. In T. Crane (Ed.), *The contents of experience* (pp. 196–215). Cambridge University Press.
- Matilal, B. K. (1986). *Perception: An essay on classical indian theories of knowledge*. Oxford University Press.
- Matthen, M. (2012). How to be sure: Sensory exploration and empirical certainty. *Philosophy and Phenomenological Research*, 88(1), 38–69. <https://doi.org/10.1111/j.1933-1592.2011.00548.x>
- Matthen, M., & Cohen, J. (2019). Many Molyneux questions. *Australasian Journal of Philosophy*, 98(1), 47–63. <https://doi.org/10.1080/00048402.2019.1603246>
- Phillips, S., & Vaidya, A. (2024). Epistemology in classical indian philosophy. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Spring 2024). <https://plato.stanford.edu/archives/spr2024/entries/epistemology-india/>; Metaphysics Research Lab, Stanford University.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549. <https://doi.org/10.1111/0029-4624.00277>
- Richardson, L. (2014). Space, time and Molyneux's question. *Ratio*, 27(4), 483–505. <https://doi.org/10.1111/rati.12081>
- Russell, B. (1911). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society*, 11, 108–128. <https://doi.org/10.1093/aristotelian/11.1.108>
- Schwenkler, J. (2012). On the matching of seen and felt shape by newly sighted subjects. *I-Perception*, 3(3), 186–188. <https://doi.org/10.1068%2Fi0525ic>
- Young, B. D. (2020). Perceiving smellscapes. *Pacific Philosophical Quarterly*, 101(2), 203–223. <https://doi.org/10.1111/papq.12309>

Matthen, M., & Cohen, J. (2024). Molyneux's question about perceptual knowledge. *Philosophy and the Mind Sciences*, 5. <https://doi.org/10.33735/phimisci.2024.11749>



©The author(s). <https://philosophymindscience.org> ISSN: 2699-0369

Open Access

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

