9

Do You *See* Space? How to Recover the Visible and Tangible Reality of Space (Without Space)

Jenann Ismael

Abstract

In the search for a theory of quantum gravity, there are strong theoretical pressures that have pushed in the direction of theories in which space (or spacetime) is not present at the fundamental level. The task of recovering the appearances is especially pressing in such theories. This chapter looks at the cognitive processes that produce spatial experience to better understand the empirical constraints on such theories. There is no question that we have immediate awareness of the visible and tangible reality of space, but what that awareness amounts to, and whether it supports the requirement that space has to be recovered as concrete external structure, is not something that has received enough attention. This chapter fills that gap.

> If one asks what . . . is characteristic of the world of ideas of physics, one is first of all struck by the following: the concepts of physics relate to a real outside world, that is, ideas are established relating to things such as bodies, fields, etc., which claim 'real existence' that is independent of the perceiving subject-ideas which on the other hand, have been brought into as secure a relationship as possible with the sense-data. It is further characteristic of these physical objects that they are thought of as arranged in a spacetime continuum.
>
> (Einstein, Letter to Max Born, 5 April 1948 (section II, p. 170))

9.1 Introduction

When you open your eyes and look around, you see an evolving spatially structured world of objects that can be heard, felt, smelled, approached, and grasped.

Jenann Ismael, *Do You See Space? How to Recover the Visible and Tangible Reality of Space (Without Space)* In: *Philosophy Beyond Spacetime: Implications from Quantum Gravity.* Edited by: Christian Wüthrich, Baptiste Le Bihan, Nick Huggett, Oxford University Press. DOI: 10.1093/oso/9780198844143.003.0009

There are strong theoretical pressures coming from the search for a theory of quantum gravity to move to theories in which spacetime is not present at the fundamental level.[1] Recovering the appearances in such a theory is non-trivial, and has been largely presumed to have to take the form of showing how space (or spatiotemporal structures) emerges (at least approximately, in the appropriate limit) as concrete structures, prior to the introduction of agents.[2] Very little explicit attention has been given in this setting to spatial perception. My chapter tries to fill that gap.

The first part looks at how the brain coordinates visual, tactual, and sensorimotor information to stabilize a conception of space as the common arena of both perception and action, and looks at what it means to *see space*, or to have immediate experience of the world as spatially structured. The second part draws some lessons for attempts to recover space and time in research in the foundations of quantum gravity.

The point of the chapter is to begin to bring the process that produces spatial experience into view, so that we can better understand what the empirical constraints are for a theory in which spacetime does not appear at the fundamental level. Readers who want to cut to the lessons for quantum gravity research might start with the second part and work backwards.

9.1.1 Quantum Gravity

The problem of quantum gravity is the problem of obtaining a complete quantum description of gravitational phenomena. Gravity has been inextricably linked to the geometry of spacetime by General Relativity, and any quantum theory of gravity has to preserve that link. There are two main classes of approach: one that results from quantizing a classical theory of gravity[3] (i.e. General Relativity), and one in which spacetime and gravity are both emergent from something else.[4] The main problem for both programs is the empirical challenge of recovering classical space(-time) in the appropriate limit. This is a straightforward matter of empirical accuracy, since the world is manifestly spatially structured.[5]

[1] Those pressures are in evidence throughout this volume but see in addition the references in note 5.

[2] In practice, it is not actually easy to say what it is for spacetime to emerge because it is not clear which features of spacetime as it appears in classical physics we want to regard as necessary for a structure being called 'spacetime'. But it is presumed that it is at least essential that the full emergence of a classical spacetime does not require the introduction of agents. It is perhaps worth adding that we remain neutral on the question of whether space is discrete or continuous, as well as on the question of whether space is a substance or a network of relations.

[3] See Teitelboim (1982), Kiefer and Weber (2005), and Oriti (2009).

[4] Oriti (2014), Padmanabhan (2015).

[5] See Huggett and Wüthrich (2013a,b, 2018), Huggett (2017, 2018), and Wüthrich (2017, 2019).

The usual way of thinking of the architecture of the universe in the emergentist paradigm is that we start from the familiar world of everyday sense, and peel back the layers revealing the physical structures that underlie the macroscopic surface of the world. Lower levels generate structures that scaffold higher levels, and as we descend to lower levels the structures become increasingly austere and increasingly unfamiliar. Colors and sounds get peeled away. The temporal asymmetries arise at an intermediate level, as a result of asymmetries in the distribution of matter. When the material contents of the universe are stripped away, what remains is a four-dimensional Lorentzian manifold in which space and time are united in a way now familiar. Below that level, we find lattices, or causal nets, or whatever one's favored theory offers for fundamental structures.

The presumption in this way of thinking is that we should be able to recover a concrete realization of general relativistic spacetime at a quite deep level, long before there is any thought of introducing agents. This program has encountered difficulties, both technical and conceptual, and there is some interest in reexamining the assumptions that are limiting how people are approaching the problem. Learning about how spatial perception works can make us less instinctively wedded to this way of thinking. In the first part of the chapter, I'm going to shed light on what it means to have the kind of immediate awareness that we have of space. In the second part, I'm going to draw some lessons for research in quantum gravity. The move toward ontologies in which space is not found at the fundamental level is one that shakes concepts at the very foundations of our thinking. The point of the chapter is not to provide definitive answers, or to push us away from any of the ongoing efforts to recover space in familiar ways. It is rather to introduce a little more sophistication about spatial perception into the discussion of quantum gravity in the interests of opening up other avenues of approach.

9.2 Part 1

9.2.1 Mach on the Construction of Phenomenological Space

When you open your eyes and look around, what you see is a three-dimensional space, external to you, in which objects distinct from your sensations are arranged.[6] What we are finding when we lift the veil and look at what is going on below the threshold of conscious awareness is that what is *given* to you in perception is the prepared product of a lot of processing by the brain. The brain

[6] Sometimes, you lose your bearings, and your experience becomes a whirl of confusing patterns of sounds and light. The difference between that disorienting whirl of sensation and the orderly ordinary world of coherent world if there, is not in general a difference in the information coming in through the sensory pathways, but in what the brain does to organize that information.

is working quite hard to integrate neural signals coming in through different sensory pathways into a unitary framework. In cognitive terms, space *is* that unitary framework. It provides a single frame of reference in which the objects of sight, sound, touch, and smell are jointly located, allowing us to share information across those pathways and use that information to guide movement.

Questions about the role of space in experience have a philosophical history running from Berkeley and Hume, through Kant, Strawson, and Evans.[7] Although it is not customary nowadays, there was a time when physicists read philosophers and thought that it was part of their job to give an account of the natural world that included the observer and accounted for aspects of human experience that are closely tied to interaction with the physical world. Einstein was engaged with philosophical thinking throughout his life. And Mach wrote a series of essays for *The Monist*[8] in which he discusses the nature, origin, and development of our concepts of space from points of view of sensory psychology, history, and physics. Although not right in all the details, Mach's essays were remarkably prescient.

There were three essays: I. On Physiological, As Distinguished from Geometrical, Space, II. On the Psychology and Natural Development of Geometry, and III. Space and Geometry from the Point of View of Physical Inquiry. In the first essay he discusses the difference between the 'physiological spaces' associated with the senses and what he calls 'geometric space'. By 'geometric space', he means what we would think of as physical space: an external dimension that contains the material bodies that are objects of visual, tactual, and auditory sensation. In the first essay, Mach argued that physical space was derived from 'physiological space', which was itself built out of atomic information delivered by the senses and motor activity.[9] The discussion proceeds as follows. Mach observes that animals have sense organs composed of one- or two-dimensional arrays of sensory elements (photoreceptors in the retina, for example, or tactile receptors in the skin). He says that such an arrangement is well suited to providing information of two kinds: (i) information carried by the response of a given element (a given tactile receptor, for example, might be geared to indicate pressure or temperature), and (ii) information provided by which element in the array is responding. The first tells us something about the condition of the thing we are getting information about; the second tells us where that thing is located. Mach's conclusion is that

[7] Berkeley (1988), Evans (1985), Guyer and Woods (1996), Strawson (1959).

[8] *The Monist* is an International Philosophical Quarterly. The papers came out between 1901 and 1903 and were published as *Space and Geometry* in 1906.

[9] There's some unclarity on a number of points in Mach's text. I've interpolated rather freely, since I'm using him here mostly as an expository tool. I apologize for the scholarly violence.

> [T]he perfect biological adaptation of large groups of connected elementary organs among one another is thus very distinctly expressed in the perception of space. (Mach 1906, 13)

According to Mach, an array of elementary sensory receptors makes available a 'physiological space' of a certain character—visual space, tactile space, auditory space, etc. About these he says:

> The physiological spaces of the different senses embrace in general physical domains which are only in part coincident. Almost the entire surface of the skin is accessible to the sense of touch, but only a part of it is visible. On the other hand, the sense of sight, as a telescopic sense, extends in general very much farther physically. . . . Yet, loosely connected as the different space sensations of the different senses may originally have been, they have still entered into connection through association . . . (Mach 1906, 15)

The physiological spaces associated with the senses are spaces in a purely formal sense. They consist of a set of elements with a relation defined on them. The fact that they form a space in the formal sense is no reason yet to think of them as having anything to tell us about physical space. To avoid confusion, I will refer to Mach's 'physiological spaces' as 'sensory manifolds', reserving 'Space' (capitalized) as a name for physical space. The pitches to which the human auditory apparatus is sensitive, for example, can be arranged along a one-dimensional continuum, from lower to higher. Receptors (hair cells) in the cochlea are arranged as a one-dimensional line of hairs sensitive to different pitches. Although vision presents itself as inherently Spatial, there is no more intrinsic reason to think of the sensory information delivered from the two-dimensional sheet of the retinae as Spatial than there is to interpret different pitches as indicating different locations in physical space. Nor is there any reason in general to think of the sensory manifolds associated with sight, touch, smell, and sound to have anything to do with one another.

What Mach suggests is that the sensory manifolds acquire a Spatial significance at the same time they acquire a significance for one another. They acquire that significance in virtue of being 'coordinated by association'. The associations in question are established through motor engagement. A creature that gets sensory information through different pathways and *acts* in the space that it gets sensory information about can establish robust, counterfactual-supporting relationships between the contents of different streams. It can use perceptual feedback on internally initiated movements to test relationships between what it sees and hears and feels.

The important upshot for our purposes is that Spatial import is not intrinsic; it is *acquired*. And it is acquired by coordination of the low-dimensional manifolds associated with the senses. Once the manifolds are coordinated with one another, a creature knows (immediately and non-inferentially and in a way that expresses itself in perception-guided movement) how to approach or retreat from things represented in its visual field. A dog knows how to get to the place from which the smell of meat is coming; a deer knows the direction in which to run to get away from the sound of danger. If you or I see a cup on a table in front of us, we know how to reach out and grab it, and that the steam and the temperature mean it will be hot against the lips. All of these implications are *extrinsic to the visual array*, but they are embodied in the Spatial significance that array holds for us. They are products of the way that that array is coordinated with other sensory and motor manifolds. And they are part of what it means to see something *as* located at a particular location in physical space.

The idea that sensory states acquire spatial significance by being coordinated in a manner that brings with it a battery of sensorimotor skills is supported by much of what cognitive science has taught us about spatial perception. There is room for dispute about the interpretation of particular experiments, and arguments to be had about the philosophical significance, but the overwhelming evidence is pointing clearly to a connection between Spatial phenomenology and motor engagement.[10]

It is tempting to think of this as though the creature already has a notion of physical space (an internal map, or a representation of some kind)[11] and the sensory manifolds acquire spatial significance by being associated with locations in *it*. But that is not the suggestion here. The suggestion is that the notion of physical space is the result of coordinating sensory manifolds, where coordination is a purely mechanical operation that precedes representation. Here's a little child's game we might use to illustrate. Suppose that you are given a bunch of patterns scribbled on papers of different sizes. You are not told anything about what they have to do with one another, but you find that there are regions of overlap that fit together in such a way that there is a way to piece them together by superimposing them onto one another, so that they form a coherent image of a three-dimensional object. And now you get the idea that what the patterns are is glimpses—from different angles—of the three-dimensional object rendered in the image. When you find parts of the pattern that seem to represent the same point, you put pins

[10] For an excellent summary of the evidence, see Briscoe and Grush (2017). The article also contains a nice survey of philosophical disputes about its significance. The reference list there is up-to-date and almost complete.

[11] I am using 'notion' in a way that is meant to span the space between the most rudimentary, pre-conceptual perceptual content and a fully articulated concept.

through them, in some cases overlaying multiple images, and utilizing the extra dimension you need to get all of the overlapping parts to fit together.

Once the construction is in place, the low-dimensional images appear as partial and redundant glimpses of the three-dimensional object, and they have a collective logic that is simpler than the sum of their parts. When I say they have a collective logic that is simpler than the sum of their parts, I mean that they have a collective geometry that is simpler than the conjunction of the patterns, if they are treated as independent. The power of this simplification expresses itself in the physics that would result if the object was put into motion. There is a drastic reduction of degrees of freedom in the assembled object that makes its physics much simpler than that obtained by treating each point on the diagrams as an independent degree of freedom. Once the object has been constructed—i.e., once you have 'solved for the object'—it would be natural and efficient for you to *see the patterns through the filter of this construction*. You would look at individual patterns for what they told you about the object, retaining information about the object, throwing away information about the patterns. You'd exploit the fixed connections between images embodied in the geometry of the object to bring information from one image to bear on another.

9.2.2 Coordination

Something like this process of construction, or triangulation, is going on in the mind as it weaves patterns across sensory manifolds into a coherent image of a three-dimensional world.[12] It is more complicated in a number of ways. It is not simply a matter of superimposing one image on another, but of weaving together streams of information about a world that is itself in flux. Identification of co-representing points is replaced with coordination of different sensory manifolds.[13] Although the contents are overlapping, the relationships among contents of streams are themselves changing (so, for example, if you have your hand on an unmoving object, and you turn your head, your visual array will change even though your tactual experience will stay the same; avert your eyes, turn your head, move your legs, and you disturb the connections between sensory arrays). The task

[12] There is a temptation to understand the kind of mapping in question in representational terms: we have a little map of egocentric space and the physiological spaces and we have little arrows representing mappings of the various sensory manifolds into the map of egocentric space. The problem with that is that—before any of this has any significance—the physiological spaces need to be coordinated with one another in a purely mechanical way. The physiological spaces are lined up, or coordinated, with one another in a purely mechanical way. Once they are coordinated with one another, egocentric space is a high-level virtual manifold stabilized out of coordination between the sensory manifolds that is made explicit in the form of mental images and concrete representations of space.

[13] Simplicity can be measured here in overall degrees of freedom.

for the mind is to piece all of this together into a coherent image of a single world, introducing as many degrees of freedom as it needs in order to accommodate variation over the streams individually, but utilizing regularities across streams to simplify the overall construction.

Just to give some sense of the difficulty of the task, consider a simplified example from Grush:

> Imagine a creature with two eyes, each of which can move in the creature's head. The creature has an arm with shoulder and elbow joints, which can bring the hand within the visual purview of the creature. There are a great many manifolds at work here. There is a 2-D manifold for each retina, a 2-D manifold for the orientation of each eye in the creature's head. There is a 2-D manifold for the orientation of the creature's head with respect to its body, and a 3-D manifold for the position of the creature's hand with respect to the torso (2 degrees of shoulder freedom, and one of elbow freedom). Let us also suppose that there's some bright object in front of the creature, which it can sense visually, as well as via touch. The project of c-coordinating the region available to vision and the region available to tactile sensation demands some sophisticated orchestration. Before they can be c-coordinated, some region of one must be identified with some region of the other. The only regions that will be stably identical are regions that are in a space centered on the creature's torso. All other regions will slide past and through each other as the creature moves this or that body part (for instance, as I move my head, my craniotopic space slides through my torso-centered space). So each must be stabilized to a reference frame anchored to the torso, as follows. The bright object will project an image onto each of the creature's retinae. However, the location of the object relative to the creature's head cannot be determined on the basis of the representations in the retinal manifolds, because the eyes may move while the head and object remain stationary, and this will change the position of the image on each retina. But, if given access both to the retinal images and to the orientation of each of the eyes in the head, the position of the object relative to the head can be fixed. That is, by appropriately s-coordinating the retinal and eye position manifolds, one can stabilize a 3-dimensional wedge anchored to the head. This 3-dimensional wedge is, as a set of ordered elements, a new manifold.[14] (Grush 2000, 68–69)

[14] See also Grush (2009, 311–345). Grush distinguishes 's-coordination' from 'c-coordination'. s-coordination involves simple identification of elements across manifolds. c-coordination involves recognition of additional degrees of dynamical freedom in the relationship between elements. I have used 'coordination' generically to cover both cases.

We're not yet finished. As Grush notes:

> This visual wedge is of little immediate use, however, because the head is free to move—something can remain at the same location in the wedge, while changing location relative to the torso and arms, if the head moves. This visual wedge needs to be stabilized with respect to the torso. This is done through stabilization with respect to a 2-dimensional manifold that provides information about the orientation of the creature's head relative to its torso, information provided by the muscles which control the position of the neck, mechanoreceptors which give information about the head's orientation, efferent copies of commands to move the head, and vestibular information. When this visual scene wedge is s-coordinated with the position of the head with respect to the body, then it becomes possible to stabilize a visual region with respect to the creature's torso. That is, for every pair (a, b) where a is an element in the visual wedge, and b is a specification of the orientation of the head with respect to the body, there is an associated point, an element in a manifold anchored on the torso. The region accessible to touch will be stabilized in an analogous manner, only though the s-coordination of rather different manifolds—limb joint angles, positions of stimulation on the skin surface, etc. When these two higher-order manifolds, visual and tactile, have been stabilized with respect to the creature's torso, it will then be possible to establish stable regions of overlap. These regions of overlap then underwrite the c-coordination of the higher order manifolds stabilized with respect to the creature's torso. The c-coordination of the 'visual space' and the 'tactile space' allows for the understanding that the felt thing and the seen thing seen are the same thing. (Grush 2000, 74)

The example illustrates how the physical act of coordinating sensory manifolds results in a higher-level virtual space, centered on the organism's body. This space is a three-dimensional manifold, anchored to the torso, which can serve as a common frame for all the lower-order (sensory as well as motor) manifolds that are coordinated with it. The result of the process is a collection of sensory manifolds coordinated and assembled on a platform that allows motion relative to one another, and that is itself moving through space. A physicist will think of this in terms of establishing relations between points in coordinate representations that are moving with respect to one another, and will think of the process of construction as a generalized best-matching procedure, where the only thing that is given is the manifolds, and the topology (including dimensionality) of the embedding space is up for grabs.

And all of this suggests, again, that neither sensation nor action carry Spatial content intrinsically. The notion of Space, as a three-dimensional manifold that is the common arena of perception and action, is literally built by the coordination of sensorimotor manifolds. It arises first as a higher-level virtual space, represented only implicitly in the stabilization of the relations between lower-level manifolds, made explicit later in maps and mental images of ourselves as embodied agents in space. Not all sensory manifolds have Spatial significance, because only some are coordinated with resulting higher-level manifolds. So, for example, retinal locations have Spatial significance, while pitches and temperatures (ordinarily) do not. Creatures must coordinate sensation and action through sensory systems and motor systems that are all connected to the torso on a variety of movable platforms (neck, shoulders, etc.) and a higher-order space centered on that common platform.[15] This platform is what we think of intuitively as our Point of View. The name is not entirely apt, because it is not really a point and it is not really a view. It is a frame of reference, with as many dimensions as the space it represents and all of the asymmetry needed to keep things aligned. And it is not tied to vision in any essential sense. What serves as the Point of View here is really the body as a whole.

Although the resulting construction sounds very complex if you try to think of it in terms of coordination among sensory surfaces, as soon as you realize that you can organize all of this complex information about the relationships between sensory and motor manifolds by just thinking in terms of information coming through sensory surfaces distributed over a body moving through Space, things snap neatly into place. Everything that we said about the complex network of evolving relationships among sensory streams is encoded in the view of yourself as an embodied agent seeing/hearing/touching/smelling a three-dimensional space through the sensors of a body whose movements you control.

We can recover all of these relationships by thinking from the outside in, or we can consider the inverse problem that one would face working from the inside out. The point of starting from the inside and working our way out is to illustrate that constructing (a notion of) space is a hard problem. The inverse problem—retrieving the relationships among manifolds that it encodes—is an easy one. That is what makes that image a good user interface. Its role, in cognitive terms, is to explicitly represent the degrees of freedom that the part of the mind that manages sensorimotor connections needs to keep track of in order to get the body moving in the right way in response to the right stimuli. It does that in the simplest way it can, eliminating redundancy and suppressing detail not needed for the task. It is also, not incidentally, the representational space in which the sorts of higher-level management tasks that require slow, deliberate thinking (decision making and long-term planning) work. The part of the mind that performs those tasks is

[15] Notice that the account of what it means for our experience to have spatial content here does not appeal to topographic maps. All of this is to be understood functionally, without appealing to neural implementation.

a very small part of the mind, from a cognitive science perspective, but the only one we know first-hand. It is the conscious part of the mind.

Once the image of physical space has been constructed, it is not tied exclusively to any particular sensory modality. It doesn't matter whether you learned where an object is by sound or touch; you know where to look if you want to see it. Sensory information represented in a Spatial format is connected immediately and non-inferentially in a web of conditional expectations for visual, tactual, kinesthetic, and auditory experience. The knowledge built into that web of connections expresses itself as practical know-how; most notably (but not exclusively) the ability to use visual information to guide motor responses. A creature that sees the world through the lenses of this construction will know immediately and non-inferentially how to approach an object seen in the visual field, reach out and grab it, avoid an obstacle, and so on.

The connection between spatial phenomenology and sensorimotor ability is developed in action-oriented theories of perception. An early version is found in Gareth Evans:

> The subject hears the sound as coming from such and such a position, but how is this position to be specified? We envisage specifications like this: he hears the sound up, or down, to the right or to the left, in front or behind, or over there. It is clear that these terms are egocentric terms: they involve the specification of the position of the sound in relation to the observer's own body. But these egocentric terms derive their meaning from their (complicated) connections with the actions of the subject . . .
>
> Auditory input, or rather the complex property of auditory input which codes the direction of the sound, acquires a spatial content for an organism by being linked with behavioral output . . . (Evans 1985, 384–385)

The suggestion here is not to reduce perceptual Spatial content to a 'muscular vocabulary'. Nor is it to reduce the Spatial content of some perceptual modalities to that of one or more others (as there was for Berkeley, for example, who looked to reduce the spatial content of vision to that of touch and movement). It is rather that it is that what it means for a creature to (for example) see space is that visual information is connected (for that creature) in a whole network of perceptual inputs and behavioral outputs. As Evans says:

> Egocentric spatial terms are the terms in which the content of our spatial experiences would be formulated, and those in which our immediate behavioral plans would be expressed. This duality is no coincidence: an egocentric space can exist only for an animal in which a complex network of connections exists between perceptual input and behavioral output. A perceptual input . . . cannot have a spatial significance for an organism except in so far as it has a place in such a complex network of input-output connections. (Evans 1982, 154)

9.2.3 Phenomenology

One's own experience (particularly visual and tactual experience) seems so intrinsically and inextricably spatial that it is hard to separate what is *in* the sensory array from what is '*acquired by association*' (to use Mach's phrase). To make this phenomenologically plausible, it helps to look at examples of experience that has been artificially enhanced in a way that allows it to acquire spatial content.[16]

Consider, for example, a device, called a Sonic Guide, which is meant to provide blind people with a sense of distance so that they can navigate by sound. It comes with a gadget worn on the head that transmits a continuous high-frequency (inaudible) probe tone and picks up echoes of that tone with a stereophonic microphone. The objects in the user's vicinity reflect components of the probe sound, in ways that depend on their size, distance, orientation, and surfaces.

The guide transforms echoes into audible sound profiles relayed to the user through earphones. The details don't matter, but just to give a sense of how it works:

1. Echoes from a distance are translated into higher pitches than echoes from nearby (as a surface approaches a user, the sound it reflects will be translated into a tone that lowers as it gets closer).
2. Weak echoes are translated into lower volumes (as an object approaches, the wearer will hear a tone that increases in volume and gets lower in pitch. An object that grows but stays at a distance will get louder, without a change in pitch).
3. Echoes from soft surfaces (grass or fur) are translated into fuzzier tones, while reflections from smooth surfaces (glass or concrete) are translated into purer tones.
4. The left-right position of the reflecting surface is translated into different arrival times of the translated sound at each ear.

As Heil (1987, 238) says, the Sonic Guide[17]

> . . . taps a wealth of auditory information ordinarily unavailable to human beings, information that overlaps in interesting ways with that afforded by vision. Spatial relationships, motions, shapes, and sizes of objects at a distance from the observer

[16] Here, and throughout, we are assuming that sensory states have a phenomenology. The question is not what gives them a phenomenology, but what makes it *Spatial*.

[17] A similar device is the Tactile-Visual Sensory Substitution (TVSS). This is a device that doesn't use auditory information but comes with an array of tactile stimulators (small vibrators), worn by a blind subject on the stomach or back, and driven by a video camera, usually worn on the subject's head. Each vibrator in the grid is driven, or not, depending on the brightness of the video camera pixel to which it corresponds. Once they've become accustomed to navigating with the TVSS, subjects report ceasing to be aware of vibrations on one's skin and becoming aware, rather, of objects and events scanned by the camera. See Bach-y-Rita (1972).

are detectable, in the usual case, only visually. The sonic guide provides a systematic and reliable means of hearing such things.

People who have been using the device for a while and are competent with it report perceiving the objects in their environment directly. Although they are aware of pitches and volumes, there is no conscious inference or reasoning out what the environment must be like on the basis of pitches and volumes. The auditory experience of someone using the device from birth, who has never known space in any other way, has a spatial content not different in kind from the spatial content that your visual states have for you. If asked what he hears, he would say that he hears an object of such and such a size located over there, just as when you are asked to report what you see, you say you see a man at the door or a book on the table. Although the sound has a quality that he can focus on and describe, the significance that the sound has for him comes from its role in his sensorimotor economy and is described in terms of what it tells him about the locations of things in space—just as although the visual array has a quality for us that we can focus on and describe, the significance it has for us comes from what it tells us about the location of things around us in space. We experience it as immediately and non-inferentially informative about the locations of things around us in space. Upon hearing a sound that represents (for him) an object located three feet from him, to the left, he would be prepared to immediately (and non-inferentially) orient toward it, point at it, or reach out to grab it, just as *seeing* an object three feet from *you* prepares you to do those things. If you or I put on the Sonic Guide and hear the same sounds, they would be just a pattern of sound signals, carrying no Spatial significance for us at all. There would be no question of using the thing to navigate, no chance of reaching out and grabbing an object at (for example) 24dB at middle C.[18] Gamers will find all of

[18] Something similar can be said about his visual experience, if he suddenly gained sight. At first, his visual experience would have no Spatial significance. It would take some time for that whirl of light and color to acquire become coordinated with the other sensory and motor manifolds. Once it did become coordinated, so that he could reach out and grab an item in his visual field, and use visual information to navigate, his visual experience would have as much Spatial content for him as ours does for us. Of course, in saying this, I'm taking an interpretational stance. One might accommodate all of the experimental data and deny (despite subjects' reports) that the Spatial phenomenology is immediate and non-inferential. Likewise, one might insist that with respect to visual phenomenology, the Spatial content is at least to some degree innate. See Briscoe and Grush (2017) for discussion of the evidence. I stand by the interpretation, but the philosophical point doesn't hinge on it. Whether the Spatial content is hard-wired or inferred, the philosophical point is that it is causally mediated by physical processes between skin and skull, and a physical theory can reproduce the phenomenology without presuming it mirrors external structure. There has to be some correspondence, but the nature of the correspondence is dictated by the role of visual space in coordinating sensorimotor behavior. Does having spatial phenomenology *consist* in possessing this pattern of sensorimotor skills? Is it merely *produced* by that pattern? I leave that question unsettled, because I genuinely don't know how to answer it, and (fortunately) I don't think we need to settle it for our purposes. Note also that one might deny that our Spatial experience is as unified as I've supposed. For example, one might deny that auditory and

this very natural. They routinely report ceasing to be aware of their own bodies as they acquire sensorimotor skills in alien environments. They aren't aware of the movements of fingers on a keyboard or hands on a joystick. They are aware, rather, of jumping or flying or propelling a virtual body through a virtual space. In some cases that virtual space is three-dimensional, but it can also be four-, five-, or six-dimensional. And it is not physically realized (as an external structure in which the gamer and game are somehow contained): it is merely represented in the two-dimensional expanse of the screen. The way to describe this is that the gaming context creates an immersed environment in which visual feedback on the muscular sensations that we ordinarily associate with finger and hand movements becomes associated with movements of the virtual body on the screen. In that environment, it is experienced as movements of the body they control, in the space in which it lives. Sensorimotor connections are locally and artificially hijacked into the virtual world.

The suggestion of all of this is that the Spatial content of our own experience is acquired by being caught up in a coordinated web of sensorimotor manifolds. It is the coordination of these manifolds that allows them to inform one another and also guide action. The difference between you and I, on one hand, and the person who has learned to use the Sonic Guide, on the other, is that his auditory sensations have gotten caught up in a web of sensorimotor connections in virtue of being coordinated with the other manifolds. This is what 'learning to use the device' amount to. At the same time that his auditory experiences have acquired a spatial content, his spatial concepts have acquired a sound profile, just as an auditory experience (and more generally, the elements in a sensory manifold) acquires a Spatial significance by being associated with locations in the higher-level space stabilized out of the coordinations among manifolds. And the elements in the higher-level manifold get their content from the entire web of sensory and motor manifolds that go into its construction.

That higher-level manifold is made explicit in the form of a mental image that can serve as input to conscious processing and is made explicit externally as when we make maps and the like.[19] We learn to picture space in our heads in part by looking at pictures and maps. By the time we are in a position to raise the question

visual space has the kind of seamless integration I've supposed. Again, here, these controversies leave the philosophical point intact. The reason for presuming that the familiar three-dimensional space of everyday sense had to be recovered as an external structure is that we seemed to have immediate awareness of it as such. If space is not even phenomenologically unified, that weakens even the prima facie case against the presumption.

[19] Older space-representing technologies—like maps—tend to represent space in a frame-independent manner, a manner that is invariant under transformations of point of view. But many new dynamical digital technologies (like GPS) represent space in terms that are relativized to (and vary with) the user's location. What gives these representational technologies Spatial content is (plausibly) the way in which they are caught up in the sensorimotor behaviors of their users.

of which comes first, we are well past the stage at which they can be separated. As we are educated, our Spatial concepts become more abstract, systematic, and articulated. How much of that additional learning filters into the way that space is presented in perception (in cognitive science lingo, how penetrable perceptual contents are to theoretical knowledge) is disputed. What is agreed on all sides is that perceptual contents directly represent things as having locations in a Spatial frame of reference anchored to our bodies. That is the sense—and I am suggesting, the only sense—in which we see space. It is the sense—and the only sense—in which we have immediate awareness of the visible and tangible *reality* of space.

9.2.4 The Upshot

There are a great many unanswered questions about spatial perception, i.e. the process that transforms sensory signals into spatially structured representations of the world. The point of the above is to give some sense of the rich body of scientific research that is bringing that process into view. The fact that we have (phenomenologically) immediate awareness of the visible and tangible reality of space makes us instinctive realists about space. From an untutored point of view, space is the dimension that houses objects that exist outside of our own minds; the things that we see, touch, hear, and smell. We are as certain of its existence as we are of the external world.

From the side-on perspective adopted above, however, the presumption that our perceptual states give us a direct, unfiltered view of reality as it is appears rather naïve. Perception is an embodied process. The products of perception were selected for their role in the organism. That role is not to provide a mirror of reality, but to provide a user interface for a part of the mind that manages sensorimotor connections. Whether, and in what sense, reality itself is spatially structured is the question of what the world has to be like to produce that pattern of sensorimotor connections. That is the question that physics has to address. The burden that the spatial appearances place on a theory is to reproduce that pattern of connections.

In the remainder of this chapter, I will look at how we might bring these lessons to bear on research on quantum gravity.

9.3 Part 2

By and large, research in quantum gravity has been shaped by a presumption that the task of recovering space demands finding space(-time) as a concrete emergent structure in the world, prior to the introduction of agents. There are various reasons for this presumption. One is that physics tends toward conservatism in its methods. Physicists are familiar with methods that involve recovering a theory in some limit, with approximations. It is not a bad policy to start with what you know and explore new methods when that doesn't work. A less benign reason is that there is a strong disciplinary resistance to introducing any talk of observation or experience into physics.

9.3.1 Maudlin's Warning: Stay Away from the Observer

Often passive and quiet, that resistance is explicit in an important paper by Tim Maudlin that has gotten a good deal of (deserved) attention from those working in quantum foundations and quantum gravity.[20] Maudlin suggests that 'naïve realism' about space might be a *sine qua non* of physics. What he says is that a theory that does not contain local beables—concrete material structures like measuring instruments occupying a public arena whose properties could serve as touchstones for our joint experiences, and whose observable properties could provide the evidence for our theories—faces a big challenge connecting evidence to theory. The argument goes like this:

> Because of the mediating place of local beables, classical physics could be tested without mention of the mind-body problem or the problem of connecting claims about experience with physical descriptions: the evidence, after all, was stated not in the language of experience but in the language of local physical facts (e.g. that the rocks hit the ground together). If the local beables and the locations are removed from the physical ontology, it is hard to see how evidential contact with the world is to be made *except* at the level of conscious experience.
>
> (Maudlin 2007, 3160)

Maudlin's worries are directed at a theory that contains no spacetime, even at the level of emergent structure.[21] The reasoning is that if there is no spacetime, there are no local beables. That means we are forced to describe the evidence in phenomenological terms, and that leads into a philosophical quagmire. Classical physics didn't have to deal with those problems because it contained local beables, and it was precisely by avoiding those questions that it made progress. He writes:

> Some connection is supposed to be recognizable between that physical description and our conscious *experiences*, even though we have nothing like a precise vocabulary with which those experiences, as such, can be described. [The]

[20] Maudlin (2007, 3151).

[21] See also Lam and Esfeld (2013), and for a response de Haro and de Regt (2018).

> mediating role of local beables obviated all these problems: there is no doubt a question about how, when we look at a rock with a certain shape, a conscious experience of certain sort arises, but classical physics could put that question off for another day (which perhaps would never come). All classical physics needs is the belief that experiences *as of* a rock of a certain shape typically *are* experiences of a rock with that shape, and the physics could take care of the rock. It is hard to see even how to begin if the physics has, in its own terms, to take care of the *experiences*. (Maudlin 2007, 3160–3161)

Later in the paper he goes on the offensive and puts the challenge in terms of a threat that Huggett and Wüthrich have characterized aptly as involving a form of empirical incoherence.[22] The worry is that if our evidence takes the form of observations of the values of local beables, and the theory entails that there are no local beables, then it undermines the very evidence that was used to establish it. The theory becomes unstable in the sense that one can't simultaneously assert the theory and hold that it is supported by the evidence. We can distinguish here a strong position and a weak position. The strong position is that the existence of local beables is a non-negotiable constraint on ontology, on pain of empirical incoherence. The weak position is that if there are no local beables, we have lost a connection between the evidence and the theory, and the only way to recover the link is to put the observer back into the picture. Maudlin actually defends the weaker position, but he wraps it in such passioned excoriation against the introduction of talk of experience or observation into physics—with gestures at the failure of positivism, and the mess of the mind-body problem—that it is clear that he thinks one would have to be a fool to defend a theory with no local beables.

It is hard to tell how seriously Maudlin intended the reference to the mind-body problem, but it is worth addressing because the whole paper is pervaded by a sense—widely shared by physicists—that whenever we start talking about experience, we are somehow in the province of a philosophical morass. One of the nicest things that has happened in the philosophical discussion of consciousness in the last twenty years is the separation of philosophical and scientific questions about the mind. A lot of the credit goes to David Chalmers who separated the mind-body problem into Hard and Easy Problems. Easy Problems, in Chalmers' terminology, are explanations of things like the capacity to discriminate, categorize, and react to stimuli from the environment; the ability of a system to access its own internal states, and to report them if queried; and the deliberate control of behavior. These are all abilities that can be functionalized and studied scientifically. The Hard Problem, in Chalmers' terms, is to give an account of the nature of

[22] Huggett and Wüthrich (2013b).

phenomenal consciousness. What makes the Hard Problem hard, according to Chalmers, is that it demands an explicit constitutive account in physical terms of what it takes for some state or mental event to be conscious, and there is a collection of arguments that seem to show that whatever one offers in the way of a functional characterization of our own conscious states will fail as a constitutive account. Being conscious is not (according to Chalmers) a functionally characterizable property.

Chalmers' primary interest was in protecting the philosophical problem about the nature of consciousness from scientific encroachment. In doing so, however, he freed the sciences of the mind to study perception, cognition, and other mental phenomena, without getting hung up on worrying about what consciousness is. The scientist can make conscious mental events an object of study, identifying them by a functional definition (e.g. introspective accessibility, availability for global broadcast),[23] and examining their internal role in human cognitive architecture and their causal connections to events in the environment. So long as she is not offering a constitutive account of what consciousness is, she is perfectly within her rights to ignore the hard problem by saying that insofar as consciousness matters to the cognitive/behavioral economy of an organism, it can be functionalized and studied scientifically. Insofar as it can't be functionalized and defined in physical terms, it doesn't make a difference to the cognitive/behavioral economy, and we can disregard it.[24]

What it means for physics is the possibility of reincorporating the observer explicitly as a part of the physical system that is implicated in the production of the experience.

9.3.2 Putting the Observer Back In

Observer avoidance might have been a good methodological strategy when the status of space was secure, and the scientific study of perception wasn't very developed. But in a setting in which physics is exploring ontologies in which space doesn't appear at the fundamental level, and we have a large and growing body of scientific research on how spatial perception works, it makes no sense. Like it or

[23] Introspective accessibility is the most natural way of identifying conscious states from a pre-theoretical point of view, but global broadcast, availability to deliberative processes, and other criteria have been proposed in the context of particular functional accounts. To say that the scientist identifies conscious events by an extensional definition (e.g. here, in terms of functional role) is to say that she has some criterion that picks them out by a definite description, in the way that we might pick out sugar as 'the dominant ingredient in lollipops' without offering that as a constitutive account of what sugar is.

[24] This is not to endorse Chalmers' dualism, but simply to allow us to study the cognitive/behavioral economy of the system without trying to answer the problem of consciousness in Chalmers' own terms.

not, the ultimate source of information about the world for each of us is perceptual experience. It is unavoidably a part of the physical chain that connects ontology to evidence and pushing it out of view doesn't make it go away.

Nor is there a need to push it out of view. What happens between skin and skull are physical processes. It became clear quite early on in the scientific study of the mind that it was most natural to speak of those processes not in neural terms, but in an information-theoretic vocabulary. Cognitive science, which models the mind as a system that processes and stores information, bringing it to bear on the determination of behavior, was born and developed much more quickly than it would have if we had to wait on an understanding of the underlying neural properties.[25] There has been an enormous amount of progress in the last half-century in studying perception in this way. In cognitive scientific terms, perception is modeled as a physical process that transforms signals coming in through sensory pathways into introspectively accessible states that represent the world as being a certain way. The process that produces those states is enormously complex and happens mostly below the threshold of introspective awareness. It can be studied scientifically using a battery of methods from behavioral experiments, to brain imaging and computational modeling. Bringing it into view makes it clear that the choices that Maudlin seems to present us with—either take perceptual experience at face value or fall into empirical incoherence—are not the only ones available. The considerations from Part 1 were intended to turn the tables on Maudlin's style of realism, by suggesting (at least implicitly) that taking experience *at face value*—if that means assuming that our perceptual states mirror the world—seems a rather naïve expectation. A cognitive scientific look at perception, which pays attention to how perceptual states are produced and the role they play in the organism, reveals that experience is neither revelation nor illusion. It does not produce representations that simply mirror the environment, but it does carry information about the world.

The way to exploit that research in the service of theories of quantum gravity is to switch from asking 'How much structure does a theory have to have in order to recover *physical space* (i.e. to recover spacetime as a concrete structure in the external environment)?' to asking 'How much structure does a theory have to have to recover *phenomenological space* (i.e. to recover the spatial content of human experience, or if you like, the spatial *appearances*)?'. We model the observer as an information-processing system, coupled to the world through sensory pathways, capable of initiating action through motor pathways, and getting sensory feedback on its actions. And we ask: what would the world have to be like to produce

[25] That is fortunate for our purposes, because it allows us describe observation without using spatiotemporal vocabulary. We need only attribute to the world a minimal causal structure: enough to support both the sensory signals coming in and the information-processing itself.

the whole structure of relationships between sensory and motor manifolds that get organized by the mind into the image of ourselves as agents living in a three-dimensional space? So long as a theory reproduces the relations between these manifolds, phenomenological space will take care of itself.[26]

Every proposal for a theory of quantum gravity will have to spatially experience its own terms. One way is to recover space as an external structure. But another is to do it without reifying space. The suggestion that we might is a natural pairing with the causal set approach, for example, and it implements the relational strategy associated most prominently with Carlo Rovelli. The relational strategy, as Oriti says so nicely, here:

> takes on board the main lesson of GR [general relativity], and it rephrases it in a way that immediately suggests a tentative solution: there is no time and no space, but only physical (imperfect) clocks and rods. The strategy amounts to identifying internal degrees of freedom of the complete system composed of metric and matter fields that can be used as approximate rods and clocks to parameterize the spatial relations and temporal evolution of the remaining degrees of freedom.[27] (Oriti, this volume)

All that we need to do to adapt this to the present setting is replace 'clocks and rods' with 'information-processing agents'.

Ultimately what physics has to recover are the spatial appearances, and that this can be done without recovering space as an external structure (and also without lapsing into idealism) is an important observation. Is it likely that this will help research in quantum gravity? Is it likely to be necessary to go all the way to the observer to recover the spatial appearances, or is space an external structure? Since the 'loss of space' is a quantum phenomenon, it concerns only structure at the very small scale and is washed out rather quickly by decoherence. So there are good reasons to expect space to emerge directly before the introduction of agents. It is important nevertheless to have the possibilities clearly on the table, to be clear on what empirical adequacy demands, and most importantly to establish that where exactly the line between internal and external structure lies is an *empirical*

[26] The pattern of relationships between sensory and motor manifolds is structure in the signal, and it carries information about the world. The job of a theory of quantum gravity is to say how the world produces that signal. The manifest image of the world—the image that the mind prepares and presents in your experience—is the message that your mind extracts from the signal. The recommendation to quantum gravity theorists is to ignore the message: focus on recovering the signal.

[27] Note also that the relational strategy has an answer to Maudlin's challenge that it is not enough that a theory allows for the emergence of structure isomorphic to spacetime. There must be something internal to the theory that makes that structure salient. In this setting, the structure has salience for an observer coupled to it in the right way. It has relational salience, not absolute salience. See also Rovelli (1991, 297–316); Gambini and Porto (2001).

question, a question of *physics*. If space is recovered as an external structure, it is because the dynamical laws deliver it in a natural way, not because we can't do physics without space.

So ultimately, it is a philosophical point that is being made here. We want to separate questions *for physics* from 'presuppositions of physical inquiry', or 'conditions of the possibility of physics', or anything like that. From a philosophical point of view, it is worthwhile to break down the idea that realism about physics demands a particularly naïve form of realism about space. Seeing in detail what is going on in the mind to produce spatial appearances tells us something about what space ultimately is for us, and about what the visible and tangible reality of space really amounts to. There is an intuitive and psychologically basic association between space and externality. It is an association elevated to an analytic connection in much contemporary metaphysics where being 'physically real' is often identified as being 'located in space'. The material in Part 1 gives us some insight into why the association is there, while also showing us why physics need not respect it.[28]

9.4 Conclusion

The idea that the mind weaves a conception of the external world out of regularities in experience is an old one in philosophy. In the Western tradition Berkeley's *A New Theory of Vision* is the *locus classicus*, and it was clearly in Mach's mind when he was writing the essays in *Space and Geometry*. Although that idea is associated with idealism, there is nothing essentially idealistic about it. Any regime that treats perception as a physical process has to recognize that, when information about the world crosses the threshold of the body, it has been transformed by the separation into different sensory streams, and the mind faces a rather complex task in trying to produce an explicit model of (the causal structure of) its source.

This has been a suggestive introduction to research that shows us how it actually does that that holds two upshots for research in quantum gravity. The immediate lesson is that it raises the possibility of a program that tries to recover spatiotemporal experience without reifying space, by letting the mind supply some of the structure that organizes sensory information into a spatial format. The more amorphous, but more important, lesson is that perhaps it is time to bring the observer back into physics. Abner Shimony saw it as one of the central tasks of a completed science to bring metaphysics and epistemology together into a mutually supporting package; the fundamental ontology of our theory should allow us to recover our experience, and our experience in its turn should be a genuine source of information about the world. He emphasized, however, that the information that experience carries about the world is itself a matter for scientific

[28] The Kantian connections are not lost on me.

illumination.[29] And he saw that when physics reaches maturity, the observer will have to appear explicitly as part of its subject matter. Only then will the circle be closed.[30]

References

Bach-y-Rita, P. (1972). *Brain Mechanisms in Sensory Substitution*. New York: Academic Press.

Berkeley, G. (1988). *Principles of Human Knowledge and Three Dialogues*. London: Penguin UK.

Briscoe, R. and R. Grush (2017). Action-based theories of perception. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 ed.). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2017/entries/action-perception/.

de Haro, S. and H. W. de Regt (2018). A precipice below which lies absurdity? Theories without a spacetime and scientific understanding. *Synthese*, 1–29. https://doi.org/10.1007/s11229-018-1874-9.

Einstein, A. and Born, M. (2004). *The Born-Einstein Letters*. London: Palgrave Macmillan.

Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.

Evans, G. (1985). Things without the mind. In *Collected Papers*, pp. 249–290. Oxford: Oxford University Press.

Gambini, R. and R. A. Porto (2001). Relational time in generally covariant quantum systems: Four models. *Physical Review D 63*(10), 105014.

Grush, R. (2000). Self, world and space: The meaning and mechanisms of ego- and allocentric spatial representation. *Brain and Mind 1*(1), 59–92.

Grush, R. (2009). Space, time and objects. In J. Bickle (Ed.), *The Oxford Handbook of Philosophy and Neuroscience*, pp. 311–345. New York: Oxford University Press.

Guyer, P. and A. Woods (1996). *The Cambridge Edition of the Works of Immanuel Kant*. Cambridge: Cambridge University Press.

Heil, J. (1987). The Molyneux question. *Journal for the Theory of Social Behaviour 17*(3), 227–241.

Huggett, N. (2017). Target space ≠ space. *Studies in History and Philosophy of Modern Physics 59*, 81–88.

[29] Shimony (1993). There is nothing certain about it, of course. It is all science, which means that it is experiment and evidence-driven conjecture. Right out of the gate we give away the prospect of the foundationalist dream of building our view of the world by logically sound inferences on epistemically secure foundations.

[30] I'd like to thank Carlo Rovelli and David Albert for comments on an early draft. I owe them both enormously. Many thanks also to Nick Huggett and Christian Wüthrich and all the participants in the Chicago area summer school on the emergence of spacetime in quantum gravity. I owe a special debt to Rick Grush, who taught me to think about space in a completely new way, and whose work I rely on heavily here.

Huggett, N. (2018). Spacetime 'emergence'. In E. Knox and A. Wilson (Eds.), *Routledge Companion to Philosophy of Physics*. London: Routledge. https://www.routledge.com/The-Routledge-Companion-to-Philosophy-of-Physics/Knox-Wilson/p/book/9781138653078.

Huggett, N. and C. Wüthrich (2013a). The emergence of spacetime in quantum theories of gravity. *Studies in History and Philosophy of Modern Physics 44*(3), 273–275.

Huggett, N. and C. Wüthrich (2013b). Emergent spacetime and empirical (in)coherence. *Studies in History and Philosophy of Modern Physics 44*(3), 276–285.

Huggett, N. and C. Wüthrich (2018). The (a)temporal emergence of spacetime. *Philosophy of Science 85*(5), 1190–1203.

Kiefer, C. and C. Weber (2005). On the interaction of mesoscopic quantum systems with gravity. *Annalen der Physik 14*(4), 253–278.

Lam, V. and M. Esfeld (2013). A dilemma for the emergence of spacetime in canonical quantum gravity. *Studies in History and Philosophy of Modern Physics 44*(3), 286–293.

Mach, E. (1906). *Space and Geometry in the Light of Physiological, Psychological and Physical Inquiry*. Chicago, IL: Open Court Publishing Company.

Maudlin, T. W. (2007). Completeness, supervenience and ontology. *Journal of Physics A: Mathematical and Theoretical 40*(12), 3151.

Oriti, D. (Ed.) (2009). *Approaches to Quantum Gravity: Toward a New Understanding of Space, Time and Matter*. Cambridge: Cambridge University Press.

Oriti, D. (2014). Disappearance and emergence of space and time in quantum gravity. *Studies in History and Philosophy of Modern Physics 46*, 186–199.

Oriti, D. (2021). Levels of spacetime emergence in quantum gravity. In C. Wüthrich, B. Le Bihan, and N. Huggett (Eds.), *Philosophy Beyond Spacetime*. Oxford: Oxford University Press. This volume.

Padmanabhan, T. (2015). Emergent gravity paradigm: Recent progress. *Modern Physics Letters A 30*(03n04), 1540007.

Rovelli, C. (1991). What is observable in classical and quantum gravity? *Classical and Quantum Gravity 8*(2), 297–316.

Shimony, A. (1993). *The Search for a Naturalistic World View*, Volume 1. Cambridge: Cambridge University Press.

Strawson, P. F. (1959). *Individuals. An Essay in Descriptive Metaphysics*. London: Routledge.

Teitelboim, C. (1982). Quantum mechanics of the gravitational field. *Physical Review D 25*(12), 3159.

Wüthrich, C. (2017). Quantum gravity from general relativity. In E. Knox and A. Wilson (Eds.), *Routledge Companion to Philosophy of Physics*. London: Routledge.

Wüthrich, C. (2019). The emergence of space and time. In S. Gibb, R. F. Hendry, and T. Lancaster (Eds.), *The Routledge Handbook of Emergence*, pp. 315–326. London: Routledge.