The Problem of Cognitive Domains © 2019, 2025 by Dan Bruiger

> "Mortals suppose that the gods are born and have clothes and voices and shapes like their own. But if oxen, horses and lions had hands or could paint with their hands and fashion works as men do, horses would paint horse-like images of gods and oxen ox-like ones, and each would fashion bodies like their own." — Xenophanes

1. The domain of consciousness

Notions of objectivity or absolute truth imply that there is some way the world really is, apart from anyone's actual perceiving of it. This, paradoxically, is the appearance the world would have if there were no observers for it to appear to. It is naive to think that the universe simply *is* the way it appears to human observers, just as it is to think it appears the same to all possible observers.

Kant drew a distinction between the realm of appearances (*phenomena*) and the realm of what intrinsically exists (*noumena*, or "the world in itself"). According to this distinction, the former is by definition what consciousness has access to, while the latter is inaccessible. An observer attempting to give shape in imagination and thought to the noumenal realm can only draw upon the accessible phenomenal realm. Though it is *not* the reality itself, for evolutionary reasons the realm of experience and thought must at least *reflect* what is real, insofar as it must permit the creature's existence. The phenomenal realm serves to *map* the noumenal territory. The unique dilemma is that we can conceive that territory only as it is portrayed to us in our map. I refer to this as the *problem of cognitive domains*, because it involves conflating the domain of the representation with the domain of what it represents. This is a category mistake, since there are essential qualitative differences between map and territory. The territory is literal while the map is symbolic. There are also quantitative differences, since a map is a highly selective version of the territory. Such reduction leaves out an indefinite amount of information.

A *domain* is the set of elements upon which some operation is to be performed, such as a mathematical function or mapping. A *cognitive* domain is some defined level of information processing—a data set that can be operated upon to become the input for further processing. As a mathematical concept, a domain is a product of human definition. The concept of domain is an element in the domain of concepts; but the world-in-itself cannot be an element of that domain. It is only with this caveat that we can speak at all of the domain of the noumenon or the world-in-itself.

The problem of cognitive domains arises when the domain that is the output of a process is recycled as its own input. This occurs, for example, when the physical world that appears in consciousness must be presupposed in order to explain the construction of this very appearance in consciousness. The phenomenal world we know in consciousness is an output of the brain, along with motor behavior. However, we can only conceive its input in the terms of this output. The world as we perceive and conceive it is taken as the causal basis and input for the brain processes that derive our perception and conception of the world, with its causal processes. There is no other access to the world-in-itself than those representations that appear in consciousness and thought, in response to the world-in-itself.

This dilemma concerns every form of cognition, including scientific theorizing; for, a mind has only the realm of its own representations (whether perceptual or conceptual) in which to speculate about possible other realms, such as an objective world or an absolute truth. Physics is an alternative form of cognition, with inputs from instruments rather than natural senses. In ordinary cognition, by default we take the world of appearances for what the world must be in itself (naïve realism). In scientific cognition, what is often taken for the world-in-itself is by default the physicist's constructed world. The biologically-based enterprise of cognition thus leads up to and includes the physicist's version of reality, which is then recycled as the starting point for that enterprise! Schopenhauer likened this bootstrap operation to the Baron von Munchausen's magical feat of lifting himself out of the water by his own pate, thus saving rider and horse alike from drowning!

This circularity characterizes the epistemic situation of the brain sealed within the skull, which may be likened to that of a submarine navigating by sonar or a windowless aircraft flown by instrument. The navigator or pilot has only the input of electronic signals from remote sensors (not a video feed) from which to plot a spatial map and a safe course. An essential aspect of such metaphors is that there is no direct access to the outside. The pilot or navigator cannot see outside, but must *infer* what lies "out there." Unlike the navigator or pilot in the metaphor, however, the literal brain has *never* been outside the confines of the skull. The challenge of the brain is to create that experience, of an external world, in the first place. On the basis of its map, the brain must invent its notion—indeed, its perception—of the territory. Furthermore, it must experience the map *as* a real external world. Effectively, it creates a virtual reality that it deals with as though it were genuine reality. This projection of the map as the territory is a normal adaptive function of biological brains. Naïve realism is the natural state for an organism that lacks a concept of map and territory, in which the world just transparently appears as real and external. It is only problematic for self-conscious beings, who have such concepts, and for whom the map is thus self-referential.

2. Science as cognition

The problem posed by consciousness is hardly the only problem of cognitive domains, nor the only instance of circularity in scientific thinking. For example, in reasoning backward to the big bang, one may wonder at the meaning of time in an environment where no time-keeping (cyclical) processes could yet exist. We use categories formed in the present era to try to grasp an era or conditions for which such categories may be meaningless, just as we use categories appropriate to one scale that may not work on another.

Science is our culture's official cognitive organ. Like the submarine navigator, the scientist's task is to create a map of the physical world from theorizing about instrument readings. The scientist has the familiar domain of experience to draw upon while knowing it to be a construct of the brain. The task for the scientist's brain is then to form a further construct, in a different cognitive domain, which empowers the human observer in ways that ordinary experience does not. Yet, the purpose of both maps is to facilitate survival. The scientific map extends the range of sensory input with instruments and expands the territory with theory. In this way, the scientific map becomes an ersatz territory to replace the natural one presented by the senses.

The entities that populate the scientific map play a somewhat different role than the ordinary objects of perception and may behave differently. In the microscopic realm, for example, "particles" do not behave like ordinary things, and do not even seem to be individually identifiable. On the largest scale, the universe as a whole can scarcely be treated as an object in an ordinary sense. Yet, natural things are ambiguous even on the human scale. (How well defined is a tree?) The physicist overcomes this natural ambiguity by explicitly *defining* the relevant entities to begin with. The entities of physics, like those of mathematics, are not natural things but theoretical constructs used to account for observed phenomena. Yet these artifacts are often treated as literal realities that (circularly) are supposed to give rise to the phenomena concerned. There is an essential, categorical difference between precisely-defined theoretical constructs and natural things, which are inherently ambiguous. Because the model is well-defined, it may be mistakenly assumed that the phenomenon it models is equally well-defined. Moreover, there is a tendency to deal only with simple phenomena, idealized in such a way that can be successfully treated mathematically (especially with differential equations). To promote theoretical models as actual elements of nature may lend them an imprimatur they do not merit and may encourage the idea that a final and complete theory is possible. In other words, scientific cognition is potentially closed and self-contained in a way that nature is not.

Later generations tend to take concepts literally and for granted, which earlier generations intended merely as heuristic. For example, the notion of 'field' was first a mathematical device before it was later endowed with ontological reality. The physical significance and reality of Minkowski's 4-dimensional continuum is still debated. For Dalton, the 'atom' had been no more than an accounting trick, not a real entity. In Newton's time, even the concept of 'force' had been controversial. In the 20th century, the concept of 'energy' shifted from being a property of matter to being substantial in its own right. It may be argued that history has vindicated these shifts or eventually will. That does not guarantee that such reification is justified in principle or in all instances. In our time, we see similar metaphysical transformations under way. While the concept of 'dimension' is mathematical, the number of *physical* dimensions has proliferated to include eleven or more; even a fractional number of dimensions has been proposed! 'Information' is being re-conceived as the fundamental building block of physical reality—rather than an elaborated product of the physical brains of highly communicative beings. Apparent anomalies in the behavior of gravitation, 'dark energy' and 'dark matter', are now considered to make up the bulk of *substance* in the universe. Physics is replete with such examples. Moreover, text-book and science writers tend to present current ideas as established fact. In their revisionist view, the entities recognized by the current generation of theorists are taken as the real and definitive constituents of the universe that were there all along.

Some quantum properties are at once found and made in the course of the experiment. For example, the very act of detection renders the position of the particle definite. The "spin" of an individual particle is held to be a property of an individual entity but is little more than a convenient name for a statistical effect in experiments of the Stern-Gerlach type. Bohr understood that quantum detection events cannot follow the same conventions as in the classical realm, where the thing observed is unaffected by the act of observation. Quantum properties cannot inhere strictly in the objects concerned, as they seem to in classical realism. Yet, we have only the classical realm derived from ordinary experience as a basis on which to conceive quantum reality. The quantum realm confounds ordinary expectations when we assume that notions based on familiar experience should apply there as well. While we can know which *kinds* of particles exist, a given particle cannot necessarily be identified as distinct from another. Quantum phenomena are collective and statistical in essence, yet there is a natural tendency to conceive individual micro events as though they were occurring on the macroscopic scale. The logic grows circular when properties on the macroscopic scale are imagined to emerge from events on the microscopic scale, while the latter are implicitly conceived in macroscopic terms.

3. Sensory and other domains

Every creature has its cognitive domain, whether or not we think of it as having conscious experience; for, all creatures make discriminations in some domain. Many rely on sense modes radically different from the familiar five, such as sensitivities to water pressure (depth), polarization of light, infrared and ultra-violet light, magnetic or electric fields, etc. Each sense modality defines a distinct cognitive domain. Consider the two cognitive domains of sight and sound. The sight of a starting pistol going off in the distance with a puff of smoke, and hearing the sound of the shot an instant later, are two experiences registering the same real event in two cognitive domains; neither perceptual event is held to cause the other. The visual appearance of a wound to one's body occurs in a different domain from the felt experience of pain associated with it. A distinct level of processing within the nervous system may also be considered a cognitive domain. These considerations would apply to artificial cognitive systems as well.

Each sense modality has a specific quality—what visual experience is "like" as opposed to what auditory experience is like. There may be something that can be learned through vision that cannot be learned through hearing. However, the specific *quality* of visual experience, as opposed to auditory experience, provides not knowledge about the world but about the sense modality itself.

Objective knowledge of the world is abstracted from sensory experience. In principle, it should be invariant over observers, sense modalities, or cognitive domains. It refers to *differences* in the phenomenal world but not to the phenomenal world itself. It is knowledge of *structure*, which is objective insofar as it should be the same for all human observers.

The biologist is no more exempt from the problem of cognitive domains than the physicist. As an open system, an organism is immersed in an environment with which it exchanges information as well as energy. But this is the situation as described by a human observer. The organism may not have a concept of its environment in the way that humans do, let alone concepts of information, energy, etc. It appears to human observers that creatures perceive and act upon the environment we perceive, while their internal representations of this environment are limited by their cognitive abilities and brain power. But this view is highly prejudiced. The very idea of an environment is an intrusion that imposes a human cognitive domain upon the organism. It may be that all that the organism ever deals with are transformations of its sensory surfaces (as perceived by humans), through acting on (what humans perceive as) its environment in such a way as to maintain them within certain tolerable limits. The irony of that possibility, of course, is that the human observer is in the very same boat. From that point of view, scientific theories are no more than aspects of human self-regulation!

While conscious experience is the culmination of elaborate processes, which themselves remain outside awareness, every stage in a flow of information processing may be considered a distinct domain. Each domain, in other words, is the subject of propositional knowledge about the presence of certain entities or events. Out of this knowledge emerges a new domain, which can in turn be searched for higher-level patterns.

A representation is a mapping from one domain to another. It may be a propositional representation, in the way that a book (which consists of sentences) maps or represents the subject it covers. While an image seems to be an analog representation, it too can be understood in propositional or digital terms. Whether or not a specific pixel on a computer monitor is illuminated constitutes a proposition, and from the ensemble of such propositions a global analog representation emerges that embodies the accumulated data of prior stages of information processing. This constitutes a distinct domain. We must bear in mind that it is the human brain, not the computer, which constructs the cognitive domain of the image from the domain of scintillating picture points. The (human) brain may then search the domain of the image for propositions in yet higher domains of meaning: esthetics, heroes, plots, etc. Each level partakes in a hierarchy of constructed domains. Analog and propositional are thus terms relative to the level or domain in question, with a dialectical relationship in the hierarchies of the nervous system.

A language one fluently understands constitutes a different domain than the collection of sounds ones hears as gibberish before learning the language. Similarly, the babble of the senses, before it has been processed in the nervous system, is not the same domain as the experience which results from that processing. Confusion arises around the idea that we can have some access in awareness to stages of sensory processing prior to the end product that we do in fact experience. The notion of *sense datum* was invoked by Locke and later philosophers as a kind of hypothetical entity in the Cartesian style. Sense data, presumably, are what we *would* experience if we *could* experience the domains of sensory input and other pre-processing stages of perception. Of course, the nervous system is not designed to include conscious access to these levels; consciousness is by definition an end product. One is able, however, to exercise special attitudes toward whatever appears in consciousness. Every painter, for instance, knows how to "flatten" visual space, to see objects not as three-dimensional things but as bounded shapes and areas of color. One may fancy that such objects of introspection bear some resemblance to domains of sensory pre-processing, but in fact they constitute a *further* domain as artifacts of conscious attention. To imagine that these afterthoughts are somehow the input to cognitive processing simply recycles a domain of output as its own input.

Similarly, the notion of *qualia* has been problematic in philosophy of mind because it involves confusing distinct domains. Qualia are the "raw feels" of sensory experience, such as the redness of the color red or the hurtfulness of pain. Yet, the fact that we are able to "bracket" these qualities as aspects of the phenomenal realm, rather than as aspects of the world, is a consequence of our awareness of being subjects. 'Sense-data', 'qualia', 'raw feels', etc., are artifacts of attention, when experience is bracketed as such, defining a cognitive domain distinct from that of the world.

4. The visual basis of objectivity

The world has survived the comings and goings of generations of observers. While the *existence* of the universe may not depend on the existence of observers, certainly our perception and knowledge of it does. The universe we know in experience and thought is a product of our mental activity in specifically human cognitive domains. But surely the world-in-itself must exist "out there," in its own right and not just in the cognitive domain of some creature? If cognition is a map, surely there is a territory. Yes, of course; the dilemma is how to conceive it.

One can grasp that phenomenal experience is not the noumenal world-in-itself—just as a film of a real scene is not the scene itself, but constitutes a separate domain. One can imagine the scene un-filmed, but one cannot picture the world-in-itself un-pictured. We are, as it were, trapped in the domain of film! In the same way, any notion of an objective reality outside all cognitive domains is paradoxically a subject's idea in a cognitive domain.

Observer independence is held to be an essential property of physical reality, at least in classical physics. But objectivity is not the absence of a subject, but a practice by a subject. In that practice, the scientific ideal is to eliminate observer idiosyncrasy. An "objective" portrait of the universe should be the same for all (qualified) observers; there should at least be intersubjective agreement.

We could call the god's-eye view "Newtonian," in which the role of the subject plays no part. This is a view from no-where and no-when in particular. It contrasts with its historical rival, the "Leibnizian" view from a unique perspective, looking outward in all directions toward all other points of the surrounding universe, each of which affords a potential unique perspective. Knowledge in that scheme necessarily derives from perceptions of a particular subject and depends on them. In contrast, the Newtonian outlook is reflected in the Minkowski space-time continuum, the so-called "block universe," in which events are point locations in an amalgamated space and time, as though seen from an arbitrary perspective outside it. The Leibnizian outlook considers the world from the unique point of view in space and time of a given subject. The Newtonian commitment to objectivity rejects this subjectivism, though it cannot logically avoid the notion of point of view itself.

An individual observer is indeed unique insofar as no other physically embodied observer can share precisely that point of view in space-time. But observation in physics follows an idealized standard based on the relative interchangeability of observers. What makes the standard possible is agreement to focus on what observers can reliably agree upon. This excludes the idiosyncratic or "private" aspects of experience, such as emotions and "qualia." That means eliminating any percept that cannot be quantitatively confirmed by others. The observer's *personal experience of observing* is irrelevant, lest it compromise the reliability of the observation. The field of discussion is the physical world, not the anecdotal experience of the observer.

Objectivity is usually identified with the visual sense. It is easy enough to grasp that the quality of pain does not exist in the world, or as a property of the hurtful stimulus that caused it, but is rather a judgment made by a nervous system. It is much harder to grasp that the same reasoning must apply to visual sensations. Visual qualia have no more objective existence outside a nervous system than does pain. This seems plausible in the case of color perception, for which scientific instrumentation has accordingly substituted measurement of quantities such as wavelength, frequency, intensity, etc. These quantifications seem to constitute a distinct domain from color qualia, even if they roughly correspond. But it is less clear in the case of percepts such as shape and size, where one assumes that vision reveals, more or less, the world as it really is. The *measurement* of space and time seems to take place in the same visual domain as the *perception* of space and time, if with greater accuracy; but the measurement of frequency seems to take place in a different domain than its sensory perception as color. However, the dissimilarity between these cases is more apparent than real. For, one way or another, it is *difference* that is detected, whether in sensory perception or in measurement. Structure is the underlying property revealed by appreciable differences. The sensory experience of color *is a detection of structure in the world*: namely, the fine structure of light. Among other considerations, it involves a biological estimate of frequency or wavelength. As in the case of shape and size, measurement by instrument may improve accuracy.

Classical science attempted to eliminate all subjective qualities, such as feelings, from its descriptions. In theory, this leaves only properties accessible to any observer, regardless of subjective state. In actuality, it means a reliance on the visual sense to provide a more "objective" relation to the world than the other senses. Vision is the paramount distance sense. Visual acuity lends itself to quantification. The ubiquitous presence of light favors objectification, which make consensus relatively possible. Precisely because of the distance from the stimulus, which allows time for reasoning and considered action, vision seems more detached than, for example, touch or smell. The ideal of science, modelled on the precision and the literally objectifying quality of vision, has been to describe the world as a realm beyond the limits of the senses. And instrumentation has vastly extended these limits. The ancient Greek thinkers, and their Renaissance admirers, sought to distill out of the flux of phenomena the unalterable real aspects of matter. They set about to create a theoretical domain transcending sensory experience, yet this conceptual realm was in fact based on the visual sense.

Yet, the greater lesson is not to mistake the visual world for the world-in-itself. Neither should one mistake physical concepts that are derived from the visual sense, which include position, velocity, momentum, etc. Physics may be our best shot at an objective description of the world, but objective does not mean independent of the subject. Ordinary perception evolved to map the world as it was relevant to humans in pre-history, in such a way as to permit survival in that world. The world has changed and so have human beings. But scientific cognition remains under the same constraint as natural cognition: it must at least permit survival. Through technology, science has transformed the human world and the natural environment as well. It seems to have greatly empowered the species. Yet the jury remains out concerning its ultimate adaptive value. Only after its extinction can one say in hindsight that a creature's cognition was inadequate.

5. The third person

Scientific or objective description is supposed to take place from a "third-person" point of view. Yet all description is necessarily in the first person: from the point of view of the describer. So-called objectivity may skirt acknowledgement of the describer's personal experience as such. But all observers stand in a first-person epistemic relation to the world—either through the natural sensory instrumentation of the observer's own body or through some external device that extends the observer's agency and epistemic relationship to the world. The fact that one can estimate the frequency of light either by means of one's color experience *or* by means of a spectrometer does not place these measures on a different epistemic footing. The measuring system is not "objective" because it has no viewpoint of its own. It is not the device that is making the measurement but the observer as a conscious agent!

First and third person accounts are distinct domains of description. Though temperature and pressure are first-person *bodily* experiences, they are also macro properties of molecules (for example, of a gas) that may be described quantitatively in third-person language. But what this actually entails is that the observer notes the reading on some instrument that is not his or her own body, such as a thermometer or pressure gauge. (Nevertheless, the observer is using some other natural sense modality—usually vision)—to make the measurement.) This reading is held to correlate with an "objective" property of the gas. To experience this same property "directly" is to "measure" its effects using one's own sensory organs as instruments.

A so-called third-person account isolates a system by excluding the part of it doing the observing from the part that is observed. This strategy avoids the problem of self-reference. But it also constitutes the two systems as distinct domains. If we try to expand the account to include the observer as part of the system observed, we simply redraw the line somewhere else, effectively repositioning a new observer outside the new system. So on in an infinite regress! The deep truth revealed here is that the epistemic subject can never be a phenomenal object. One can see one's body, but not one's self.

The scientific picture—at least as mis-portrayed in popular accounts—pretends to depict the noumenal face of the world before any nervous system has gotten ahold of it. Though the utility of scientific models depends on their mathematical power of prediction, their intuitive appeal refers back to ordinary experience, cloaked as objective reality. Science rejects the cognitive domain of ordinary (first-person) experience as *merely* a cognitive domain, only to construct *new* cognitive domains that it does not acknowledge as such. It holds its theoretical entities to be the building blocks of the objective world-in-itself, while these conceptions often remain tied to acts of visual imagination.

6. The problem of consciousness

The scientific problem that consciousness poses is a problem of cognitive domains. The mental is a distinct domain from the physical. Physicalism proposes to reduce the mental domain to the physical one. Paradoxically, however, the physical domain is a construct *in* the mental domain. The problem chases its own tail.

The so-called ‘hard problem of consciousness’ reflects this circularity, in view of which it is no more plausible to reduce the mental to the physical than vice-versa. Dualism is built into the self-referential nature of consciousness. The problem is not to eliminate dualism by explaining consciousness in strictly physical terms. Rather, it is to understand how an organism internally relates the two disparate domains for its own purposes: how it maps its internal representations to the putative external world. (This might apply to other epistemic agents as well, such as fully autonomous robots).

There can be no strictly physical causal account of consciousness because causality takes place within a single domain: that of the physical. By definition, physical causes only affect physical things, in a rote manner that does not account for the autonomy of organisms. Causality thus fails to relate to the other domain, which is the mental. We need a different concept, capable of relating domains to each other and of clarifying the nature of this relationship. I propose that the concept of intentionality, broadly understood, can serve this purpose. *Intending* is the internal act of an autopoietic *agent* (an organism, whether natural or artificial). Specifically, it is an act of mapping one domain to another. (In contrast, causality is a third-person attribution by an agent, observing from outside a system that is not considered an agent in its own right.) Agents *act*, both internally and externally; matter only passively reacts to other matter. In this sense, an agent *makes* connections within itself, in contrast to events within it that might be externally caused. Intending is not limited to conscious intention, but includes it. While intending is first-personal, external observers are free to theorize third-personally about the intentions of other agents (specifically the internal operations of organisms), just as they can theorize about causes.

How does this help to explain consciousness? First, it tells us to stop barking up the wrong tree. An explanation of consciousness in strictly causal terms is not possible, any more than Cézanne’s portrait of an apple is explained by the apple used for a model or by the paint used by the artist. Like that of the artist, the intentionality of the brain is the bridge between the domain of the apple and the domain of its image, whether on the retina, on canvas, or in the mind’s eye. The artist creates the image, brush stroke by brush stroke, each of which relates one domain to another.

Of course, the metaphor fails for the reason we have been exploring: there is no access to the domain of the apple-in-itself. The “painting” that the brain produces is not arbitrarily created; but neither is it a copy of anything. It has a basis in reality, but is symbolic. It is a process of mapping, but not an isomorphism. The sensory experience is produced through neural events that are internal acts within the “artist.” But these acts ultimately have to do with the self-maintenance of the organism; they relate only indirectly to the world, insofar as the organism depends on it. The “accurate” portrait of the apple is the one that enables the organism to eat it. That could be a very different portrait for the worm or the deer than for the human being.

Just as the brush strokes *create* the domain of the image on canvas, it is the mapping process itself that creates the inner domain of consciousness. I call this fundamental creative act of mapping *fiat*. Intentionality consists of it. So to speak, something in the brain says *let there be light*—and there *is* light! It may be objected that this is hardly satisfying as an explanation of how the phenomenal realm arises from neural activity. But, I suggest that such objection is based on an inappropriate demand for causal explanation. It is more fruitful to examine the organism’s relations with its environment, in order to speculate on the *reasons* for which it makes its internal connections, as opposed to physical causes that underlie those connections. It is better to take a Leibnizian approach, in which one tries to put oneself in the brain’s shoes as an agent, than a Newtonian approach in which the brain is no more than molecules jostling about. There is no effective bridge from the domain of physical causes to the domain of phenomenal experience, but there is a bridge to be built from the domain of intentions.