

## Is there a TV in my head?: Content, Functional Mapping, and the Myth of the Given

Just what are we talking about when we talk about the **content** of perception? There are a number of different questions which are often asked about the precise nature of this content, such as to what extent it is **representational**, and to what extent this representation is **conceptual** in character. However, these questions usually gloss over what it is to talk about perceptual content *qua* content, in their haste to talk about it *qua* representation or conception. On this basis, it is all too tempting to account for not only the *specific character* of the content, but also its *mode of individuation*, in **phenomenological** terms. This makes perceptual content into whatever is contained within an **introspective domain** to which we have some sort of special access. Wilfrid Sellars' attack on the myth of the given has provided us with good reason to doubt the epistemic authority such special access to our inner states purportedly provides, and thus to doubt the efficacy of any account of perceptual content that gives pride of place to introspection. This is his critique of what Jim O'Shea has called the **epistemic given**. The aim of the present paper is to shed some light on his critique of the other side of the myth, or what O'Shea calls the **categorial given**, by working through some of his other ideas. This involves demonstrating constraints governing accounts of perceptual content on the basis of the *explanatory demands* placed upon them as accounts of perception, and the *explanatory resources* available to them as accounts of content.

### 1. Perception and Explanation

To begin with the explanatory demands, it seems to me that there are two distinct explanatory enterprises that any account of perceptual content must contribute to. On the one hand, there is the **epistemological** enterprise of explaining the *general role* that perception can play in empirical justification independently of *variations* in the causal structure of perceiving agents. This means telling a story about how the **sensory inputs** fed into a causal system can gain the *normative significance* of warranting moves within the space of reasons, in a way that could apply not only to different human beings, but to stranger creatures such as aliens or artificial intelligences whose sensory capacities and overall causal economy diverge radically from our own. On the other hand, there is the **psychological** enterprise of explaining the *specific role* that perception plays in the causal economy of particular perceiving agents. This means telling a story about how the sensory inputs fed into a causal system contribute to the production of **behavioural outputs**, in such a way as to give us *predictive purchase* upon the behaviour of creatures with particular types of causal structure. Insofar as introspection is a further kind of perception, it is something that needs to be

accounted for in both stories, and this gives us reason for caution in using it to tell these stories.

Although these enterprises are distinct, they are also importantly intertwined. Any epistemology that cannot account for the way differences in the causal structure of our sensory capacities can affect their role in empirical justification will have failed to get a grip on the *causal dimension* of perception, or its connection to **sensation**, and any psychology that cannot account for the way perception can supply us with reasons that cause us to act one way rather than another will have failed to get a grip on the *normative dimension* of perception, or its connection to **rational agency**. We only have an account of perception, be it epistemological or psychological, when its causal and normative dimensions are properly connected. Any account of perceptual content as something that plays a role in both epistemological and psychological explanations must couch it in terms amenable to both of these dimensions. The first insight that we can take from Sellars here is that the proper interface between these two dimensions is the use of **functional explanation** in empirical psychology. To use a phrase he is fond of – we must understand perceptual content in terms of the way it fits into the **wiring diagram** of the perceiver.

Causal explanation in general works by applying **explanatory schema** to systems that facilitate the development of **predictions** about the way that they would behave under various possible conditions. These schema provide us with more or less *general* ways of organising **counterfactual reasoning** about these possibilities, thereby enabling us to draw *specific* conclusions about how they would behave in any given set of circumstances. A **functional schema** enables us to develop predictions by treating a system in analogy with **practical reasoning**. For instance, by allowing us to treat its **parts** as **means** in relation to the **whole** as an **end**. This lets us describe the causal role of the systems' components in terms of **success** and **failure**, and thereby to organise our counterfactual reasoning about the causal relations between them in terms of the way failure **cascades** throughout the system. The explanatory power of a functional schema thus lies precisely in its introduction of the possibility of **malfunction**. Crucially, this is an essentially normative notion, which can be leveraged in connecting the epistemological and psychological accounts of perception.

## 2. Content and Explanation

Moving on to the explanatory resources, it strikes me that the purpose of the notion of content is to talk about how two seemingly *different* states of distinct systems can nevertheless be the *same* in another sense. There's obviously a trivial sense in which this can be the case without warranting the

ascription of content at all, such as the sense in which two houses may be in a state of disrepair, even while the precise nature and extent of their disrepair may differ. But to say that they both share a *content* here would be to say nothing more than that they share a *property*, and thus entirely redundant. This is not to say that states whose similarity consists in common properties cannot share content on that very basis. Two VHS tapes that share precisely the same magnetic properties will share precisely the same informational content. The question is how we can say that an information storage medium that works on different causal principles, such as a Betamax tape or a DVD could share the same content as the VHS tapes. I'm going to use this comparison with information storage media as the guiding analogy through which to think about the explanatory role of perceptual content. I think we should aim to think about what it would be for states of distinct perceiving subjects *qua* causal systems to have the same content, in much the way that states of distinct information storage devices can have the same content, before we address what it would be for this content to be perceptual, and in what sense this makes it representational and/or conceptual.

In order to extend the individuation of the content belonging to states of causal systems beyond mere similarity of properties, we must consider relations of **isomorphism** between them. This means breaking down states into the *features* of which they are composed and the *relations* between them, and then developing a way of **mapping** these to the features and relations that compose the corresponding states. This **mapping schema** (or **morphism**) allows us to count some feature of one state as equivalent to a feature of another despite differences in their properties, insofar as it occupies the same role within the system of relations that constitute it. We can then determine if two states share the same content on the basis of some sufficient degree of correspondence between their components. This would let us see a VHS and a Betamax tape as possessing the same content insofar as there is some way of mapping the magnetic properties of one to the other, despite the differences between these properties. However, the problem with such **pure isomorphism** is that it can be arbitrarily extended in ways that undermine any possible explanatory role it could have. We can potentially construct arbitrary mappings that pair the magnetic traces on the VHS tape with price patterns in the stock market, or pair the digital encoding of the DVD with a sequence taken from the binary expansion of  $\pi$ . Insofar as it completely severs the individuation of content from any concern with the causal capacities of systems and their states, pure isomorphism precludes identity of content from playing any role in causal explanation.

We can avoid these problems of pure isomorphism by using the **functional roles** of the states and their components to constrain the mapping schema. In order for this to work the relevant states must

be **variable** features of the wiring diagram of the system, and their variations must be functionally correlated with **variable outputs** of some sort. For example, the VHS tape contains a length of material whose electromagnetic properties vary in delimited ways, which will produce suitable variations in the patterns of light emitted by a TV set to which it is appropriately connected. On this basis, it is possible to produce a **functional mapping** from the components of the one set of variable states to another in terms of **common outputs** to which they are functionally correlated. For example, we can produce a functional mapping from VHS to Betamax that maps their distinct variations in electromagnetic properties onto one another in terms of similarities in the patterns of light they engender when suitably connected to a TV set. In essence, we treat the states possessing content as isomorphic with one another insofar as their components can be mapped onto the same set of functional roles, which are themselves isomorphic with the output they are supposed to produce. There are some issues here regarding the fineness of grain of this isomorphism, but we will leave them be, in favour of two final points about functional mapping.

The first point is that we need not characterise content individuated this way as representational. It is all too tempting to say that the content *represents* the output that it is functionally correlated with. For instance, it can be tempting to suggest that two VHS tapes represent the same movie, or that two records represent the same piece of music (something Sellars is guilty of at one point). Here it is important to remember that ‘movies’ or ‘pieces of music’ are just as **abstract** as contents, and their correct showing or performance is subject to their own set of criteria. One might nevertheless think that content represents **raw output**, such as the patterns of light or sound with which the information storage media are correlated. However, we should resist the temptation to identify representation with mere functional correlation, as it arises principally because the analogical character of functional explanation invites us to treat the relevant states *as if* they were instructions for performing certain actions.

The second point is that this gives us a way of talking about **form** as well as **content**. The ways in which the components of the relevant states can vary may be classified in terms of more general functional roles they play, and these classifications provide more or less abstract forms that corresponds to the content of their specific variations. Put another way, form consists in the **functional invariants** that delimit the variations in which content consists. It is the structure of those elements of the wiring diagram in which the variability of the content bearing states are encoded.

### 3. Perception, Representation and Conception

We can now turn to considering which states of a causal system deserve to have something called perceptual content ascribed to them, and to what extent they are representational and/or conceptual. Here I am going to follow Sellars, who takes perception to be the transition from a causal input to a conceptually articulated output, or from **sensation** to **conception**. I'm also going to endorse his account of the nature of this conceptually articulated output, which is thought of as a sentence or as a state whose content is equivalent to a sentence. Conceptual content is principally understood as the functional role that a sentence or its component expressions play in the game of giving and asking for reasons, which is composed from three types of move that are made with sentences: **language-entry** moves (perception), **intra-language** moves (inference), and **language-exit** moves (action). Two sentences have the same content just insofar as the same moves can be made with them, and two component expressions have the same content just insofar as their contribution to the roles of the sentences they compose is the same. Conceptual content can be derivatively possessed by states of systems insofar as they are appropriately functionally connected to the possibility of linguistic output, i.e., of actually making moves within the game. This means that the sense in which the transition from a causal input to a conceptually contentful state counts as perceptual is derivative upon the sense in which a language-entry move is perceptual. Finally, I'm going to say that conceptual content is representational insofar as I agree with Sellars that it **signifies** things in the world, but I'm not going to go any further into the tortured question of how representation can be reconstructed out of inference here.

On this basis, I think that anything worth the name perceptual content will have to be possessed by states of the mechanisms involved in the whole process of moving from sensation to conception. Borrowing a term from Brandom, we could say that perceptual content must be possessed by some state of the **perceptual mechanism** underlying a rational agent's **reliable differential responsive dispositions** (RDRDs). However, there are potentially many candidates for this, insofar as there can be numerous subsystem states with variable functional outputs involved in the processing of sensory information into conceptual content. Restricting ourselves to our own visual systems for the moment, there is everything from the pattern of activation on the back of the retina (Quine's infamous **stimulus meaning**), through states of information processing systems in the visual cortex, to states of a global system that integrates information from various sources and makes it available to other cognitive processes (such as Thomas Metzinger's **Phenomenal World Model**). Moreover, because content bearing states compose into further states, it is equally possible to talk about

combinations of any or all of these. There are multiple layers of information processing between sensation and conception any and all of which can be the subject of functional mappings insofar as they output to other layers or to the ultimate conceptually articulated product.

The question is now what it would be to say that any of these states possessed representational yet non-conceptual content. This is what Sellars calls **picturing** as opposed to signifying. The obvious thing to do here is to reach for the notion of isomorphism once more, and to say that content which is already individuated by functional mapping is representational just insofar as there is also an isomorphism between it and some state in the world, which it is thereby taken to represent. For example, the content of the VHS and Betamax tapes may be the same insofar as they record the data from the same security camera, and they may then be taken to represent the same events insofar as there is an isomorphism between this content and the relevant events. However, this suffers from the same problems with arbitrary mappings we discussed earlier, insofar as we can conjure up isomorphisms with other potentially stranger events or states. We might then suggest that this isomorphism is constrained by the **causal origin** of the relevant states. However, this would leave us saying that if the tapes were warped in precisely the same way by the same magnetic field, that they thereby represent that magnetic field. Sellars' solution to this problem is that the isomorphism must be more deeply tied into the functional role of the content bearing state.

In 'Being and Being Known', he illustrates this using the example of a robot that stores information on a similar magnetic tape, which he takes to picture its environment in virtue of an isomorphism between the state of the tape and the state of the environment. However, he also claims that:

This picturing cannot be abstracted from the mechanical and electronic processes in which the tape is caught up. The patterns on the tape do not picture the robot's environment merely by virtue of being patterns on the tape. In Wittgenstein's phrase, the 'method of projection' of the map involves the manner in which the patterns on the tape are added to, modified, and responded to by the other components of the robot. It is a map only by virtue of the physical *habitus* of the robot, i.e., by virtue of mechanical and electronic propensities which are rooted, ultimately, in its wiring diagram. (§40)

To summarise, the constraints upon an isomorphism that allows us to produce a **representational mapping** between the states of two systems does not merely concern similarity of functional output, but a more complex relation of **projection** in which the state mediates between functionally

specified inputs and outputs. For example, we can understand a variable state of the wiring diagram of a bee as representing the path to a source of nectar insofar as it varies appropriately when the bee discovers the source of nectar and produces the behaviour of returning to that source of nectar with other bees that it has danced for. This is something that Ruth Millikan has expanded upon in much more detail.

So, for a state of our perceptual mechanisms to be non-conceptual and yet representational in this sense would be for it to play a functional role in a process of systematically guiding *behaviour* in relation to sensory input that is to some extent independent of any role played by conceptual content in guiding *action*, which is derivative upon the role of sentences in language-exit moves in much the way that perception is derivative upon the role of sentences in language-entry moves. The tension here, which I do not aim to resolve, is the extent to which these states can play an active role in the move from sensation to conception without their role in guiding behaviour being similarly bound up in conception.

#### 4. Universal and Parochial Content

Having distinguished between the non-representational, non-conceptual, and conceptual contents that states of our perceptual mechanisms can bear in terms of the functional roles of their components within these mechanisms, I now want to make a further distinction between types of content on this basis, and see how this distinction can shed some light on the myth of the categorial given. The distinction I want to draw is between what I'll call **parochial** and **universal** forms of content. The former covers all forms of content whose individuation is *dependent* upon functional mappings that are specified in terms of particular causal mechanisms, whereas the latter covers all forms of content whose individuation is *entirely independent* of any particular causal mechanism. We've already made this sort of distinction in considering the difference between contents individuated in terms of functional mappings and contents individuated in terms of pure isomorphisms. The characterisation of the latter is entirely mechanism independent, but ultimately cannot play a useful role in any causal explanation, psychological or otherwise, whereas the characterisation of the former, along with the representational mappings we have discussed depends upon the causal structure of the mechanisms in terms of which the output is specified.

However, the conceptual content which is the result of perception is universal in the sense just defined. This is because the language games from which the conceptual content of states is derived

are social practices that are intrinsically **extensible**. Though any given speaker's ability to perform language-entry moves is tied to the structure of their own perceptual mechanisms, the concepts that they thereby apply are not forever tied down to those mechanisms. It is entirely possible for speakers with entirely distinct causal economies to acquire the concept as long as their linguistic usage can be appropriately mapped from within the game of giving and asking for reasons itself, insofar as the causal structure of any given perceptual mechanism is something amenable to analysis in conceptual terms. This is the same process through which we incorporate new and more accurate measurement devices into our practices, thereby bootstrapping our ability to apply empirical concepts on the back of our parochial perceptual mechanisms.

Nevertheless, this extensibility and the universality it implies does not detract from the fact that our concepts are achievements of distinct sort. Concepts display relatively fixed forms of functional invariance in much the way that forms of functionally individuated content do. This is what Sellars' means by the categorial. Form in general is the classification of functional roles. Categorial form is the classification of specifically conceptual roles. Categories are thus general classifications of concepts that constitute the general form of signification as opposed to the general classifications of causal states involve in the various parochial forms of picturing.

All this gives us a way of thinking about the illusion that engenders the myth of the categorial given. I'll quote Sellars' on this form of the myth:

To reject the Myth of the Given is to reject the idea that the categorial structure of the world – if it has a categorial structure – imposes itself on the mind as a seal imposes an image on melted wax. (§45)

I think we can reformulate the myth as follows: it consists in the idea that there is some form of *universal perceptual content* that is distinct from *empirical conceptual content* as Sellars describes it. This is just the claim that there is some specific sense in which absolutely any two sentient creatures could be said to have the same *experience* without having the same *conceptual grasp* of this experience. This amounts to the idea that I have a TV in my head that has no particular causal-functional structure. It is the postulation of a kind of content that cannot in principle be individuated, or rather, which is self-individuating. This derives from the temptation to individuate perceptual content on the basis of **what** we take it to represent (e.g., the pink ice cube we're both looking at), rather than **how** it represents it (e.g., the functional role the relevant states play in a



wider behavioural economy that incorporates the pink ice cube). This makes the posited content universal, but only insofar as it becomes parasitic upon the conceptual content of sentences we use to describe what it represents. Under the power of this myth, the appearance of some common form within introspective apprehension of this purportedly universal content thereby engenders the illusion that the world itself is impressing its categorial structure upon us directly.