LANGARA COLLEGE

Philosophy 2201 – Epistemology

Richard Johns, updated January 2024

Perception Basics

In thinking about knowledge, one of the key concepts is *belief*. The word 'belief' often is used with a religious connotation, but for philosophers a belief is simply something a person takes to be the case, something that is true from that person's point of view. For example, my beliefs include: *Langara College is in Vancouver*, *the atmosphere is mostly nitrogen*, *whales are mammals*, *Shakespeare wrote many plays*, and so on.

The concept of belief is often contrasted to that of knowledge. A person may say, for example, "I don't *believe* that whales are mammals, I *know* they are." However, if you know that whales are mammals, then obviously you take it as true that whales are mammals, so you do believe this in the philosophical sense. Knowledge is a special kind of belief, roughly a good belief, or one that has nothing wrong with it.

One kind of knowledge comes from sense perception, i.e. from seeing, hearing, touching, smelling and tasting. When you think about it, it's clear that the vast majority of our knowledge ultimately comes from this source. We know things because we or others have observed them to be so. Some philosophers (called empiricists) even claim that *all* our knowledge comes from observation and introspection. Hence, when studying knowledge, sense perception is a good place to begin. Most of our discussion will concern vision, since it is arguably our most interesting and important sense.

When we look at a situation, we gain some knowledge about it. Thus, since knowledge is a special case of belief (see above) the process of seeing results in our brain forming a belief, or a set of beliefs. When you look at a coffee mug, for example, you see *that* it is a coffee mug, that it is blue, that it is almost empty, that it is resting on the desk, next to the lamp, and so on. Furthermore, many of these beliefs appear spontaneously in our minds, without any reasoning or inference, even though other beliefs may be inferred later. For example, suppose that when you walk into the kitchen you see that a chair has been moved, so that it now stands next to the counter. Then you notice that the cookie jar, which is kept on that counter, has its lid off. Then you see cookie crumbs next to the open jar. Each of these beliefs forms immediately and effortlessly, during the observation, so that you say, "I directly saw that the lid was off". Based on this direct knowledge, you may infer that someone has been eating cookies. You reason that lid was removed to allow access to the cookies, and that the crumbs dropped when this person bit into the cookies. You also infer that your toddler was the culprit, since only such a small person would need the chair to stand on. Thus some beliefs form by a process of logical inference, based on visual beliefs, but this doesn't change the fact that vision itself is a belief-forming process – in other words, "seeing is believing".

Some more terminology. Two people can sometimes 'have the same belief', as we say, when for example you and I both believe that whales are mammals, not fish. But philosophers note that your belief and mine are technically two different things, since they have different properties. For example, your belief is something that exists in your head, while my belief exists in my head. Your belief started to exist on Tuesday last week, whereas mine didn't form until this morning, and so on. So in what sense are our beliefs the same? The sameness lies in their *content*; we both believe the same *thing*, and this 'thing' is called a proposition. You and I both believe the proposition that whales are mammals.

Propositions can of course be expressed with declarative sentences, as I'm doing right now, but a person can have a belief without forming any sentence in their mind that expresses it. When you look at a scene in front of you, for example, the visual experience that occurs in your mind is called the *visual field*. (The visual field is like a 3D movie playing in your mind. As you turn your head left, the visual field pans to the right. As you walk forward, objects at the centre of the visual field appear larger. When you close your eyes, the visual field goes dark, and so on.) The visual field is propositional, we say, since the objects in it are presented as being a certain way. Generally each object is presented as having a certain size, 3D shape and colour, and as having certain spatial relations to other objects in the field. Many objects in the visual field are also presented as being a certain *kind* of object, such as a tree, a person, a dog, etc. though ambiguous objects can also exist. Some objects in the visual field, especially people, even have particular identities. You might, for example, see your Aunt Edna walking toward you.

So let's review. When you look at something, your mind forms a visual field, which includes a collection of beliefs. These beliefs will generally count as knowledge as well, unless something is wrong with your eyes or brain. (In older literature, beliefs are also called 'judgements', by the way.) Each belief you have is an event in your mind, and thus belongs to you alone, but the thing that you believe, the proposition, is an abstract entity that others might also belief, doubt, wish to be true, reject as false, and so on. (Propositions that arise from sense experience are often called sensory propositions.) A proposition is a representation of reality, a kind of map or picture of reality that you can carry inside your mind. Of course such representations are very useful, if true, as they enable us to find our way around the world. A map that shows that actual location of buried treasure, for example, will enable us to dig in just the right place and get rich.

The notion of truth has been mentioned now and again. Truth seems to be a property of beliefs and other representations, but

what is it? When is a belief, model or map true? The standard answer is that a true belief (or map or model) is one that exactly matches or corresponds to the reality that it aims to represent.

Look at the two maps below, for example. Both aim to represent the same reality, namely part of Europe in the 16th century. The map on the left was drawn in 1522, while the one on the right was made just a few years ago. You will see that the maps approximately agree in some aspects, while diverging in others. For example they both show Scotland (Scotia) to the north of England (Anglia). But one map shows England and Scotland fused together as a single island, whereas the other represents them as separate islands.

You will likely judge the newer map to be true or correct in this regard, but what does that mean exactly? We can compare maps with each other, finding agreement and difference, and apparently we can make a similar comparison between a map and 'reality'. Truth, as understood by common sense, means agreement with reality. England and Scotland *in fact* are joined, not separated by ocean. The shapes of these countries are also not round blobs as the left-hand map shows, but have crinkly edges, similar to how they're depicted in the right-hand map.

Another useful notion, when speaking about truth, is a *state of affairs*. It's rather obvious that two very different maps can both be true. For example, one might be a map of Japan, and the other a map of Australia. How can they both agree with reality if they're not similar to each other? The obvious answer is that they each correspond to different parts of reality, or different aspects of reality. A single part or aspect of reality, that may or may not agree with a given proposition, is called a *fact* or *state of affairs*. In other words, no proposition ever aims to represent the whole of reality, and the little chunk that is represented is called a fact or state of affairs. More precisely, a fact is an *actual* state of affairs, one that actually exists. Philosophers also talk about *possible* states of affairs that aren't actual, such as Hilary Clinton being the US president in 2017. A false belief can correspond to a non-actual possible state of affairs, or (in rare cases) to no possible state of affairs at all.

Concepts, Percepts and Ideas

Let's look more closely at propositions, to see the structure within them. Propositions contain simpler elements, some of which are concepts. For example, the proposition *there is a snail on my desk* contains the concepts of snail and desk, among others. Some concepts that are commonly involved in sensory propositions (whether visual or otherwise) are geometrical concepts such as shapes, sizes and locations (e.g. *square*, *circle*, *touching*, etc.), colours (*red*, *blue*, *brown*), tastes (*sweet*, *lemony*, *salty*), smells, temperatures, textures, biological categories, types of artifact, etc. etc. Note that early modern philosophers such as Descartes, Locke and Berkeley used the term 'idea' instead of concept here.

In addition to general concepts, propositions also have components that are particular objects. A proposition might be about cats in general, such as *Cats think their owners exist to serve them*, or about a particular (individual) cat, as in *Fluffy is overweight*. Such particular individuals, that exist at least in the subjective world of

the believer if not in reality, can be called mental objects or sometimes 'internal objects'. In the case of the visual field, the particular objects within them are called *percepts*. A mental object, such as a percept, is different from concept, in that it represents a single, particular individual (e.g. Aunt Edna), whereas a concept represents a general kind or category of things (e.g. aunts).

Direct vs. Representative Realism

Suppose you're sitting in your back yard, looking at your apple tree. In that case your visual field has a particular object, namely your apple tree, filling most of it. In the previous section I called such a particular object in your visual field a 'percept'. But what really is a percept? And how is it related to the actual tree?

According to a philosophical theory called 'direct realism', the tree-percept and the tree itself are one and the same thing. In other words, when you look at the tree you are directly conscious of the tree itself. Thomas Reid, the best-known direct realist, writes for example (see the iweb Readings):

> When we see the sun or moon, we have no doubt that the very objects which we immediately see are very far distant from us, and from one another. We have not the least doubt that this is the sun and moon which God created some thousands of years ago, and which have continued to perform their revolutions in the heavens ever since.

The key word here is 'immediately', which seems to mean that seeing the sun does not involve any representation or image, of

which we are more directly aware.[1] David Hume rejects direct realism, but describes the view as follows:

> It also seems clear that when men follow this blind and powerful instinct of nature they always suppose that the very images that their senses present to them are the external objects that they perceive; it never crosses their minds that sensory images are merely representations of external objects.

A more complex alternative to direct realism is 'indirect realism', or 'representative realism'. On this view, the tree-percept is not the same as the tree itself. The tree-percept is an 'internal' object, which means that it only exists in your mind, whereas the real apple tree is 'external', i.e. it exists independently of your mind. The tree-percept on this view is a mental *representation*, or model, of the real tree.

Representative realism is obviously more complex than direct realism, in that it claims there are (in a way) two trees rather than just one. Now Ockham's Razor tells us that entities are not to be multiplied beyond necessity, so there had better be a good reason for having an extra tree.

---

[1] The reader should be warned that some contemporary views, according to which perception does involve a mental representation, are also referred to as 'direct realism'. E.g. Huemer's *Skepticism and the Veil of Perception*, 2001.

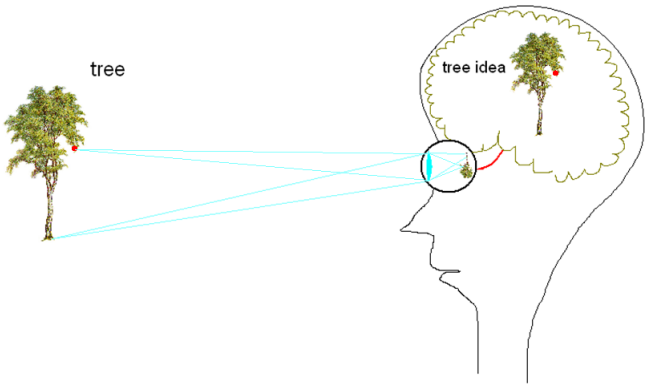

Representative realism's two (or even four!) trees

One argument for the existence of the tree percept (or idea) comes from the theory that conscious awareness is some kind of *brain* state. In that case, since we're conscious of the tree, the tree must be some kind of feature of a brain state. But the real tree, the wooden one, obviously doesn't exist in your brain, because it's out there in the yard, not inside your skull. So, any 'tree' in your conscious awareness must be a different tree altogether.

This argument is reinforced by our present scientific understanding of vision. On this theory, seeing a tree depends on light rays being emitted from various parts of the tree and entering our eyes. These rays are 'focused' by each eye lens onto the retina, which means that light rays emanating from a single point of the tree all end up at a single point on the retina. Thus some point on the retina (near the top, actually) will receive light only from the base of the trunk, whereas a different region of the retina will receive light only from a particular red apple on the tree. In this way, we say that the lens produces an *image* of the tree on the retina.

Having a tree-image on each retina is not enough to see the tree. If your optic nerves are cut, then you'll be totally blind, no matter how well your eyes are working. Some cells in the retina, called rods and cones, convert the light rays they're receiving into electrical signals, and these signals pass along the optic nerves to the brain. Somehow the brain uses these electrical signals, coming from two eyes, to construct a visual field, including a 3D model of the yard and the tree within it. Initially the retina signals are processed in two areas of the brain called the 'visual cortex', but the actual formation of the visual field is still a mystery, as far as I know. We do know for example that the colours of a percept depend on the cone cells, of which there are three different kinds, each sensitive to a different wavelength of light. If the light from an object triggers mostly your 'green' cones, then the brain will colour that object green in your visual field.

Evidence for this scientific theory comes partly from the fact that hallucinations exist, and can be as vivid and 'real' as genuine perceptions. Suppose, for example, you eat a psilocybin mushroom and consequently suffer a hallucination of an apple tree. The tree you experience, in your visual field, might seem utterly real and convincing, yet there is no external tree present. Thus, in the case of hallucination, there is a tree-percept but no tree, proving that they are not identical in this case at least.

Further evidence comes from illusions, where an object appears to be rather different from how it really is. For example, 3D movies work by showing the left and right eyes of the viewer slightly different images. When looking at a real 3D object, one that lies a few feet in front of us, our eyes form slightly different images of it, due to their different perspectives, our pupils being 2-3 inches apart. The 3D movie supplies images that correspond to ones that the eyes *would* get if there *were* an object in the space between the viewer and the screen. This fools the brain's visual system into constructing a 3D percept, in the visual field, apparently floating

before the viewer. For example, we might see a seagull fly through the movie theatre, over the heads of the people seated in front of us. This shows pretty clearly that even during normal vision, when there really is an external object present, the thing we're directly conscious of is not that external object.

The basic argument goes as follows. These examples of hallucination and illusion show that the brain has the natural capacity to construct 3D percepts that do not correspond to anything (or at least to anything three-dimensional) in the real world. Now, why would the brain have such a remarkable ability? What purpose would it serve? (Presumably its purpose is not to make acid trips or 3D movies possible!) This capacity is surely used every day, during ordinary visual perception. In other words, ordinary vision involves the brain making educated guesses about what objects exist in the space in front of the viewer, and drawing percepts in the visual field to represent those objects.

After a person comes to accept representative realism, this view might come to seem obvious, even trivial, and so direct realism seems very foolish. For this reason, direct realism is often referred to by the (derogatory) term *naive* realism. Direct realists are accused of the elementary mistake of confusing a representation with the thing represented, such as confusing the map and the territory. A person giving directions to a newcomer may, when pointing to a map, say "Here's Langara College ... here's 49th Avenue ...", but no sensible person thinks they're literally pointing at the college or the road in that case. They're just pointing at representations.

Non-Literal Representation

In philosophical terminology, a *realist* is someone who believes in external objects that (at least approximately) correspond to their beliefs, mental representations, percepts etc. However, a representative realist need not believe that a percept exactly corresponds to its external object in all respects. A map for example contains representations of various real objects such as schools, roads, rivers, parks, and so on, but the representations are not exactly similar to the real things. The representations are much smaller, for example, they're paper-thin, and often a different colour. In other respects, such as shape, they are identical to the real objects – at least approximately.

Temperatures in Canada

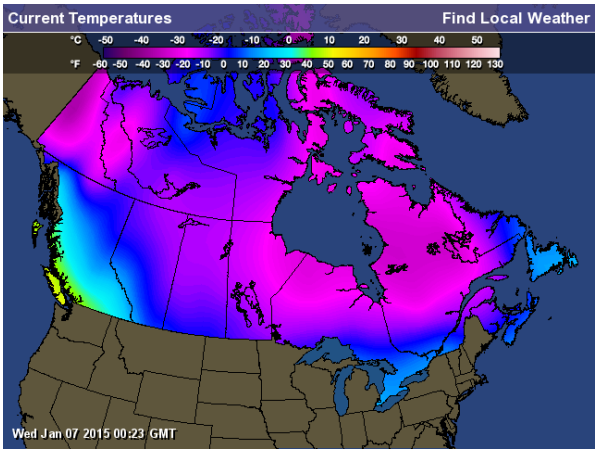

For example, the image above uses colours to represent the temperatures in various parts of Canada on a particular day. It conveys the information that south-west B.C. was fairly balmy at that time, while northern Ontario was bitterly cold. (Ha ha! Take that, Ontario!) Of course the map isn't telling us that northern

Ontario is purple, or that Vancouver Island is yellow. Those colours are merely representative, not literal. On the other hand, the shape of "Ontario" on the map is very similar to the shape of Ontario (the real, external object).

Thus, when it comes to our own percepts, in the visual field, we should be open to the notion that they may not resemble their corresponding external objects in all respects. Even if they are accurate, conveying correct information, they may do so in a non-literal way.

Such non-literal representation appears to be the case with many qualities of our percepts, including colours, smells, and tastes among others. Contemporary physics seems to have no place for anything remotely like the colours that exist in the visual field. Physics describes a world of particles in motion, and says that the perceived colour of an object depends only on the wavelength of the light that it emits. For example, longer wavelengths (around 630 nm) appear red to us, and shorter wavelengths (around 450 nm) appear blue. This is surprising to many, since there is of course no particular resemblance or similarity between the colour red and wavelengths of 630 nm, any more than there is a resemblance between triangles and the colour pink. The connection is purely arbitrary, from a physical perspective.

In other words, the purple colour of a grape is no more literal and exact than the purple colour of frigid Ontario in the weather map. It's merely representative of some other quality altogether.

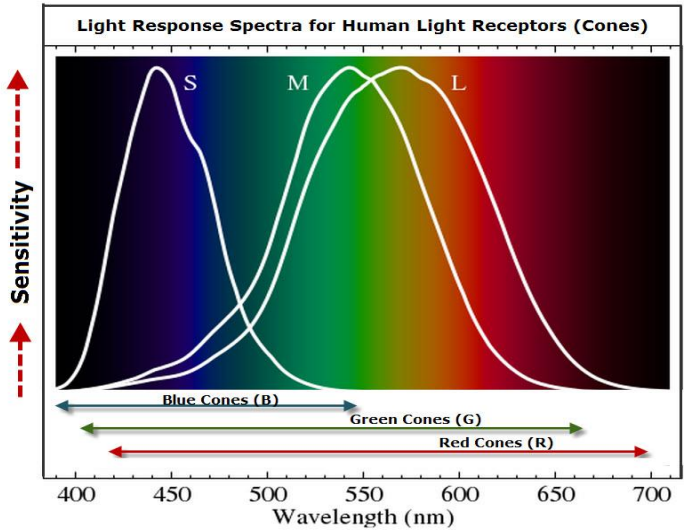

John Locke wasn't the first philosopher to propose that some qualities of percepts represent non-literally, but he invented the standard terminology to describe it. Qualities of objects that are represented literally by our percepts (or 'ideas') are called *primary qualities*. Qualities of objects that are represented by very different properties in the percept are called *secondary qualities*. Thus Locke writes,

> ... the ideas of the primary qualities of bodies resemble them, and their patterns really do exist in the bodies themselves; but the ideas produced in us by secondary qualities don't resemble them at all. There is nothing like our ideas of secondary qualities existing in the bodies themselves. All they are in the bodies is a power to produce those sensations in us."

Are there *sense data*?

It is stated above that beliefs often involve 'internal' (or 'mental') objects. For example, in order to believe that Donald Trump is the current US president, there must be such a thing as Donald Trump in my subjective world—an internal object representing The Donald must exist in my belief state. In a similar way, when I look into my back yard and see the old apple tree that grows there, this process involves the construction of an apple tree percept in my visual field.

Now, some philosophers would call this apple tree in my visual field a *sense datum* (plural: sense *data*). A sense datum is an internal or mental object that one is directly aware of during sense perception (e.g. vision). So far so good. But sense data are also regarded as having two additional features that are rather questionable:

> (i) Sense data are objects of some kind, that the mind stands in some relation to. (We are said to be *aware of* sense data, or even to *perceive* sense data.)
>
> (ii) Sense data actually have (exactly and only) the properties that appear to us. For example, when we look at a tomato, we may see it as round and red. In that case, according to the sense data theory, the tomato's sense-datum actually *is* both round and red.

(Recall that the external tomato isn't red as such, since colours are merely representative of wavelengths of light.)

At first sight it may seem obvious that sense data are the same as percepts, but these additional claims that percepts are separate entities, and that they actually have the properties that we attribute to the external object, are rather problematic.

The sense data theory talks as if perception is a relation between the conscious mind and the sense datum, as if the sense datum itself exists independently of the mind. But surely percepts are simply *components* or *aspects* of our mental states, not separate

objects that the mind stands in some relation to? If sense data are really separate objects that we perceive, then where are they? Are they physical objects? Mental objects?

The claim that sense data actually have the properties that we attribute to the external object is also hard to defend. Let's think again about seeing a round, red tomato. The tomato itself is of course round (assuming we see it right) but does the percept have to be round as well? The percept is a *representation*, and it definitely represents the tomato *as* round, but does this require the percept to also be round? Such a view seems silly in some cases. What if I see a piece of steel heated with a blowtorch, until it glows red hot, and I can feel the intense heat coming from it? Is my percept of the steel actually hot as well? If it were, wouldn't it cause burns inside my brain? If I see a football stadium, that is 150 yards long, is my sense datum also 150 yards long?

These criticisms have led many philosophers to reject the sense datum theory, and endorse instead the adverbial theory, or even some version of direct realism.

The Adverbial Theory

The adverbial theory is another form of representative realism, according to which the properties that an object appears to have, in our visual field, are not properties of the representation itself, but properties that the representation *ascribes* to the external object. For example, if we see a red, round tomato, then our visual experience *presents the tomato* as being red and round, rather than the visual experience involving a sense datum that is *itself* red and round. And in the famous case of a half-submerged stick appearing bent, there is no sense-datum that is actually bent, but just a representation that *ascribes bentness* to the stick.

To get a better sense of this difference, it can help to consider a type of representation that we are familiar with, such as a map. Suppose that a certain section of road is extremely steep, and this

fact is to be shown on a map of the area. A very literal way to do this would be to mould the map out of clay, so that its surface rises and falls in correspondence with the territory. In that case, the representation of the steep road would itself be steep! This is analogous to the sense-data theory, since the property P is depicted by a representation that itself has P. Suppose however that the map is instead to be made from flat paper. In that case one could just write 'steep' on that part of the map, or draw arrows pointing downhill. On this approach, the line on the map that represents the steep road is not *itself* steep, but it still *ascribes steepness* to the actual road. This is analogous to the adverbial theory.

A map of Dolgellau in north Wales.

Note the steep sections of road, shown by arrows pointing downhill.

Return to direct realism?

Many contemporary philosophers (e.g. John Searle, Michael Huemer, Galen Strawson, Hilary Putnam) now advocate what they call 'direct realism' as a superior alternative to both forms of representative realism. This is rather different from the direct realism that Hume dismissed, and Reid apparently supported, in that these recent philosophers accept the standard view that perception involves the construction of representations (i.e. percepts), and that perception occurs by means of these representations. How then do they differ from representative realists?

These new direct realists distinguish themselves from representative realists by claiming that during a perception, we directly form beliefs about external objects – I believe for example that (alas) *my coffee cup is empty*. I do not first form a belief about my subjective percepts (e.g. that the coffee cup percept in my visual field is empty) and then infer that the external coffee cup is empty as well. This doctrine, that perception *directly* produces beliefs about the external world and not indirectly via an earlier belief about sense experience, is supposed to be characteristic of direct realism.

In other words, these recent direct realists *attribute* to representative realists the claim that perception always occurs in two stages, where we first form a belief about our percepts, and then infer beliefs about the external world. Huemer (2001, p. 3) writes for example:

> *Indirect realists* hold, instead, that our awareness of the real world is indirect. They accept arguments like the one given above, which says that what we are immediately aware of in perception is only mental images; however, they say that we can *infer* the existence of real objects *corresponding* to our images, because that is the best explanation for why we have the sort of mental images we do.

In other words, Huemer portrays representative realism as saying that perception as rather like having a conversation with King Charles, via the King's representative. Suppose a journalist has been granted an interview with the King, but (perhaps for security reasons) the journalist is not permitted to talk to the King face to face. Instead, the journalist asks a question to the representative, who then walks into an adjoining room, presents the question to the King, receives an answer, and then returns and relays this answer to the journalist. In this scenario, the journalist is initially aware of what the representative is saying, even if he then infers that the King said the very same thing.

In response to these direct realists, BonJour agrees that perception (in normal cases) directly produces beliefs about the external world. Ordinary perceptual beliefs are things like, "My cup is empty," not things like, "I am having a sensory experience of my cup being empty." BonJour expresses this by saying that perception is *intentionally direct*. However, since views like Huemer's accept that perception involves the construction of a mental representation, and that this representation may or may not correspond to the external world (since illusions and hallucinations are possible) BonJour says that "the view that results is still fundamentally a version of representative realism".

Huemer's view that perception is intentionally direct can be illustrated by the case of watching a game on TV. When you watch a game at home like this, you're obviously not looking directly at the actual players on the field. You're looking at representations of the players, formed by coloured pixels on the screen. But someone who's 'immersed in the game', as we say, isn't consciously aware of the pixels. They're only conscious of what's happening on the field, e.g. that a goal has just been scored. The TV screen is then a kind of *vehicle* that allows them to be aware of what is happening on the field. It's possible to 'snap out of' that immersive state, and be aware of the TV itself, but it's also possible to forget about the TV and just (as we say) watch *the game*.

According to Huemer (and representative realists who hold that perception is intentionally direct) each of us is almost always in such an immersive state with respect to our visual percepts. We're not aware of the percept itself; instead the percept is a *vehicle*, or a tool, for us to be aware of the world in front of us. Representative realists sometimes make another analogy, comparing a percept to a window. You don't (usually) look *at* a window, you look *through* it, to the world that lies beyond. In a similar way, they say, you can look 'through' the TV to watch the game, and look 'through' your moon-percept to see the moon.

Why does it matter?

Representative realism is an important theory in epistemology because it is a necessary premise in an important argument for radical scepticism. Mike Huemer presents this argument in Section 2.1 "The Infinite Regress Argument", and it can be roughly summarised as:

1. Representative realism is true

2. Internalism is true

∴ I cannot know anything about the external world

Like Huemer, I think such arguments are important, even though radical scepticism is false. Since we *know* that radical scepticism is false, we can use such arguments to infer *other* facts about human knowledge. For example, Huemer concludes that representative realism is false, and others conclude that internalism is false.

I haven't yet defined 'internalism' here, but it is one of the main topics of this course. The rough idea of internalism is that a knower must *justified* in holding the belief in question, and that such justification requires that the knower be consciously aware of evidence supporting the belief. In Huemer's argument for radical scepticism, for example (in Section 2.1) the assumption of internalism shows itself in the following statement:

> "If the presence of F is to explain why I am rational (or justified) in accepting A but not rational in accepting B, it must be something that I am aware of"

Generally speaking, internalists will tend to say that perception involves a conscious inference from sense experience, as this allows perception to be 'policed' by our rational judgement.

Reason can then decide what, given all the relevant background evidence as well as the occurrent experience, is the right conclusion to draw. Opposed to internalism is (not surprisingly) *externalism*. Externalists will say that perception involves a spontaneous belief about external objects, without any conscious inference (i.e. perception is intentionally direct). Even though these perceptual beliefs do involve inference from sensory inputs, as we have seen, these inferences are carried out automatically by unconscious mechanisms.

Externalism does not allow perception to have much rational oversight, but externalists do not see such rational checking as necessarily for knowledge. 'Reliabilism', for example (the most popular kind of externalism) says that all that matters is that the perceptual process is reliable.