

What Does it Mean to Mean Something?

Introduction

In this paper I will consider broadly arguments for meaning skepticism with respect to spoken language. There are a few different flavors of these arguments, but more or less they all argue for doubting that language means what we typically think it means. In essence, these arguments put forward that language use does not imply the things we typically assume about speaker brain state. I will examine these arguments, and sketch one line of argument that can be tried to escape the rabbit hole.

The structure of the paper is as follows. I will first define meaning as a speaker's intended effect on a listener. I then will consider how to bridge the gap between minds, and what it would look like for meaning to obtain. In this context I will consider further skeptical arguments about language,

Meaning

Meaning is the intended reference of communication. When someone communicates, meaning is their intended subject matter. A speaker will have some purpose in expressing themselves, and this will be something like to conjure an image in the listener's mind. Their meaning is then that mental operation. The idea is that a speaker will intend by speaking. Since a speaker speaks to a listener, they intend to have an impact on their listener. For simplicity we will just assume throughout that there is only ever a single listener in question. So now the speaker intends to impart a message on the listener. In doing so, the speaker works with a model of the listener's mind. That is, they must imagine themselves interpreted in order to speak, since the purpose of speaking is to be interpreted, and the interpretation is in the mind of the listener. Thus the meaning content is a model of listener reaction. Therefore, in looking for meaning, we are trying to measure how accurately a speaker predicts listener reaction.

With literal speech, a speaker may be said to mean some part of the world, that part of the world that the conjured image is a depiction of. To put it simply, meaning is the part of the world a speaker wishes to draw attention to. Although in fact meaning would refer to what the speaker thinks the listener interprets the reference as, for a literal reference a speaker will plainly assume that the listener imagines the word as referring to the same thing the speaker has in mind. So in the literal case, it is reasonable to say that there is a part of the world that the speaker means. It is still the case that the desired operations are intended to take place in a listener's mind, but this just means that the speaker intends to refer to the physical object as apprehended by the listener. Vitaly, the speaker need not care about how the listener apprehends that part of the world for the reference to obtain. At any rate, in literal reference, it is possible to imagine meaning as just referring to a part of the world, even though strictly speaking meaning is still defined as that which a speaker intends to impart on a listener.

Communication

Meaning is a fairly fundamental part of language, and it is the assumption of meaning on which conversation operates. The basis of conversation is quite simply the assumption that in using language one manages to get a message across, and conversely that what comes to mind is that which the speaker intended. This makes meaning skepticism a fairly heavy blow for all of communication, since it strikes at the very foundation of language in general. Yet the question arises, how can we gauge meaning at all.

In order for communication to work, we need to have agreement across different minds. The fundamental issue is that minds are not the sort of objects that can touch each other directly. Instead, each mind exists somehow independently. Inherently, the stuff of one mind exists in isolation from that of another mind. Since minds do not directly connect, what metric, outside of language, can be used to establish agreement between minds? How can agreement between two distinct minds even be spoken of coherently?

Here we can break the problem into two parts. One section is the issue of reference, and the other is one of interpretation. Reference is meaning of a single word, while interpretation is

the internal mechanism of a sentence. Meaning exists primarily in the interpretation part of this consideration, and reference is subsumed under this header. Nonetheless, these two are somewhat distinct, as establishing reference is quite different from establishing interpretation. This is because reference is narrower, and more easily established. The question before us is then how to make sense of intended speaker effect without direct access to a speaker's mind.

Satisfaction by Agreement

First we will discuss reference agreement, and we will then leave this issue aside for some time. A possible standard for reference agreement is comparing judgement of photos. We can imagine defining a dog by looking at a series of photos judging each to be a dog or not. In this process, there may be some edge cases. Some example edge cases are: a photo with just the shadow of a dog, a photo containing only the front half of a dog, or a photo of something between a dog and a wolf. In these edge cases, one might not call the picture in question a photo of a dog. This provides a judgement of the reference "dog". Here we can try to establish co-reference by comparing our judgement with a peer. In an extreme case we can imagine looking at photos of every instance of a dog each of you has seen, from every angle and frame you've seen it. If you can agree on all your known references, it ought to be the case that you refer to the same thing in using the word "dog".

Insofar as reference consists in agreeing whether a word describes an element of the world, a metric ought to be at least theoretically possible. Since we can limit ourselves to historical instances, all a speaker and listener need to do is go through each other's lives frame by frame, agreeing on each frame whether the word in question is at play. It is reasonable to compress this into a more practical process, by looking at a large sample of edge cases, and some standard examples and non-examples for comparison. At any rate, the question of reference becomes one more or less of practicality, and less of a theoretical barrier. Since in order for a reference to obtain I really only need use agreement, and can accept that your mental image is distinct from mine.

The issue remains in the interpretation part of language. This is the meatier and much trickier issue. The question here is what is going on in the minds of speaker and listener. Here in fact we are concerned with more than agreement, and something like assent and dissent facts will not satisfy us. A way to understand this is that in talking one often cares about how their conversation partner feels. In looking for operating on things like feelings, attitudes and behavior patterns, conversation tries to get at something beyond language use agreement. Simply put, when talking one wishes to access another's mind. In order to do so, a speaker cares about how the sentence comes across. Yet without direct access to minds, meaning seems impossible to ascertain.

Speaker Models

We have now the issue of how a listener can understand a speaker. Since minds exist in isolation, how is a listener to know what effect a speaker intends to have. Sure, a listener does have a handle on the effect a speaker has, since the effect is on the mind of the listener. If a listener could indeed compare this to the intended effect, then we will have the desired metric of meaning. However, the lack of direct access to a speaker's mind throws a wrench in this process. We know that the speaker intends a certain reaction, and on the other hand we know that there is a listener reaction, yet we do not know how to connect the two, since they take place inside distinct minds.

Let us step back for a moment and recap. Meaning is the effect the speaker intends to have on a listener. A listener can gauge from their own mind the actual effect. However, a listener cannot gauge intended effect, since this exists in the speaker's mind. Since a listener cannot directly apprehend meaning, because a listener cannot know a speaker's intentions, we fail to have a way to know whether meaning works. That is, we don't have a way to compare intended effect to actual effect.

More broadly, the issue is that we cannot simultaneously examine a listener's and speaker's minds. Deeply embedded in communication is the desire to connect minds. This is the reason that meaning is all about the speaker imagining the listener's mind. Yet minds are

fundamentally separate. This separation of minds is the wrench in our program. We are seeking for a way to compare notes on language interpretation. Yet the issue at play is really how to do this, if not for within language, which is the very thing we are in doubt about.

Meeting Expectation

We can know something about the mind of another by how it responds to stimuli. If another's mind is an input-output machine, we can test its internal structure by watching it react to a variety of inputs. The input in question is the listener expectation, based on conversation up to the point in question, and the output is speaker speech, which interacts with expectation. While this ends up telling us not very much about the first person experience of the content of another's mind, it tells us a whole bunch about the structure. If we have somehow already established a high degree of reference agreement, then this sort of input-output testing gets further along the line, and we can see meaning start to take shape. The essential ingredient to move forward will be expectations.

The way out of this mess is in predicting responses. There are two parts here, the extent to which a conversation partner conforms to expectation, and the extent to which they diverge from it. Oddly enough, these both contribute to a connection of minds. When conforming to expectation, speaking in a way that can be predicted, a speaker confirms that they understand what we are saying. When defying expectation, a conversation partner lets us know their own thoughts, which were not already in our own minds. In order to do a good job at adding to a listener's mind, a speaker needs to have a good working knowledge of the listener's mental representations. Thus, expectation can be used to gauge speaker intent.

Care must be taken to ensure that we are collecting actual data. The potential pitfall in seeing confirmation in both meeting and defying expectation is quite immediate, for how could two opposites both confirm the same hypothesis?! Here the devil is in the details. The conformity to expectation is in the manner of discussion, as befitting the content. The defiance of expectation is in the content of the speech. If things are the other way around, then they indeed disconfirm meaning. If a speaker defies expectations in manner, this is a worrisome departure

which indicates misunderstanding the way things are represented in a listener's mind. On the other hand, in conforming in content, a speaker would simply mirror the thoughts of their conversation partner, and not convey any grasp of what's being discussed. These two elements used in the right way however, are the window into a speaker's mind.

It is worth stepping back here, and doing some careful bookkeeping. We need to keep track of what a listener is doing and what a speaker is doing. This is in order to ensure that our argument does not rest on some tricky slight-of-hand arising out of ignoring the speaker-listener distinction. The idea is that the speaker's interaction with listener expectation gives the listener an idea of the speaker's intended manipulation of a listener's mind. Going into a conversation, or at any rate after a bunch has been said, a listener will expect the speaker to say one thing or another. The degree to which and manner in which the speaker conforms to this expectation is taken as evidence for the speaker's intention. At the same time, the listener's experience in reacting to what's said is taken as the effect of speaking. By comparing the modeled speaker intention arising out of conformity to listener expectation with actual impact based on listener reaction, a metric for meaning is obtained. This is the degree of alignment with the listener's perception of meaning with what is spoken. Meaning here is understood as before, as the speaker's intended effect on a listener. And what is spoken is understood as the way the listener interprets what is said. At any rate, this clarifies exactly what claim I'm making, and which part of my argument goes where.

To reiterate, we are now in a position to examine meaning. We started by defining meaning as speaker's intended effect on listener's mind. We then said that determining the legitimacy of meaning would consist in comparing the speaker's intended effect to the actual effect. Yet we saw an issue in even discussing meaning, as minds are inherently isolated. We now found a way to bridge that gap. First, in pointing to reference agreement by exhaustive examples, we had a way to agree on reference. Then, by appealing to conformity to expectation, we start to have a way to get at how a conversation partner interprets us. Nonetheless, the issue of mental isolation persists, and muddies the waters of meaning. Thus we still have a lot to consider in order to establish meaning as we know it. Nonetheless, we now have enough to enter that discussion,

since bridging the mind gap has been shown to be conceptually feasible. Granted, we have got a long way to go, but we're not quite sitting and fumbling around about how any meaning is possible at all given mental isolation. We can finally move from wondering whether meaning can ever be coherently discussed to actually starting that discussion. So we now go ahead and try to get a handle on how well the above works for actually establishing agreement between minds.

From here on out I will be a bit looser in talking about meaning, and in doing so I will be practicing less technical bookkeeping. My hope is that the above discussion suffices to clarify what sort of paradigm I'm working in. This sort of looser talk will allow me to tackle higher level concepts without getting lost in the weeds. I think loose talk is appropriate for the sort of discussion I'm transitioning to, where we are gazing at the forest instead of inspecting trees. The idea is to move from what meaning is and why we have any reason to believe it might exist to why we are confident in our belief in meaning and how to navigate in dealing with meaning.

Meaning Skepticism

Meaning skepticism is fairly prevalent. There are many ways to argue for it. One recurring line of thought in meaning skepticism is the interdependence of language, since words are defined in terms of other words, there is no bottoming out, and no way to reason into meaning from the outside. The worry is how do we know that there is any sort of associated mental process whatsoever. Perhaps language is but a game that people play with each other to pass the time. Essentially the issue is that language is rich enough that it can account for use without resorting to meaning.

Another issue is in the prevalence of linguistic imprecision. Since we cannot precisely pin down speaker intent, there is inherent ambiguity. And since language is a network, this ambiguity is infectious, and casts doubt on the possibility of meaning at all. That is, if each term can mean several similar but not-quite-the-same things, higher level concepts begin to crumble. In particular, this kind of ambiguity introduces a subjectivity of language. This subjectivity would be an inherent feature of language itself, so that it could never be escaped, because there's no way

to know the content of someone else's thought, besides through language, which we saw to be imprecise and thus subjective.

Another skeptical angle to take is one in which terms fail to personify their intended references. The idea here is that words exist in relation to each other, and the fact that relations between things hold needn't imply that the relations we intend between them hold. This is a world in which we mean something by our terms, and use it to reason about the world, but our reasoning is futile. Instead of discussing features of the world as we wish, we end up discussing nothing but our language itself. To be more precise, the worry is that language underspecifies its content, leading to potential speaker disagreement. The complexity of the world relative to the simplicity of language allows for words to mean varied things.

Interconnectivity

One thing that is strikingly present in this discussion is the extent to which all of language is connected. This is a fundamental part of the above skeptical arguments. One side of this interconnectivity is present in the fact that language need rely on little outside of itself. So that an entire conversation can be held with little mention of any immediate stimulus. This speaks to the isolation of language from stimuli. This removal from immediate stimulus allows for the skeptic to contemplate the total divorce of language from any stimulus whatsoever, namely in robbing words of their referents. The problem is that meaning gains its meaning in large part by being anchored in some part of reality, but language can stand on its own in a way that is too abstract to be tied concretely to its content.

Another thing interconnectivity allows for is the permutation of the meanings of all sentences in such a way to preserve relations between words but alter large swaths of meaning. Since language is so connected, lots of sentence meanings rely on other sentence's meanings. This allows for a world in which only a few words need have their referents changed, and suddenly nothing rings true as intended. There are many highly conceptual words, and if the meaning of some of these were changed, nothing would quite mean what we think it means. That is to say, the fact that language is a network makes it sensitive to small perturbations.

Use Facts

What seems to be the operating principle in these skeptical arguments is the poverty of language use facts. The instances of language seem in the whole insufficient to determine meaning. If one is to take the totality of some collection of language use facts, as if it were a really long text taking into account all information about this collection of facts, it seems that a single meaning would not be determined. This is because a book in isolation, no matter how elaborate and detailed, cannot give itself meaning without a reader already knowing the language at hand. In studies of old dead scripts, researchers rely on chains of transmission and families of similar languages, and even so can only make the bare minimum of sense of these texts.

Use facts in general behave in this way, as a dead language. Of the many possible interpretations and referents that can be assigned to some collection of words, use facts does little to distinguish between these. It isn't that use facts don't manage to rule out some possibilities, but rather that new possibilities can be introduced at a greater rate than use facts can rule them out. In particular, Kripke has pointed out that there are infinitely many possible completions, but only finitely many use facts.

Speaker Agreement

An essential bit in the argument against skepticism lies in the way that speakers align in their assent and dissent to sentences. This high degree of agreement, in particular about basic facts, is the first thing to point to in establishing that language does indeed have meaning. The argument proceeds by saying that the most natural explanation for speaker agreement is the meaning of sentences and language user alignment. This gives belief in meaning a leg to stand on, and is essentially the reason we think language has meaning in the first place.

Skepticism from the Outside

We in practice are not skeptics, this begs the question: how in practice do we avoid the allure of skepticism? What is it that lets us know our meanings, given the poverty of use facts?

The key here lies in understanding how communication actually works. The thing we rely on is not use facts themselves, but the process in which they are embedded. It is not the use alone which determines our knowledge of meaning, but the use in conjunction with our expectations, and the way in which our intentions are responded to in conversation. This ends up being some version of speaker agreement. When speaking, our conversation partners validate our expectations, thus expressing agreement with our internal models of language, suggesting that language exists atop a platform of meaning.

Thus skepticism is understandable, while simultaneously avoidable. When looking at use alone skepticism seems entirely possible, but when we match that with our internal states and thus our expectations, the pull of skepticism is weak. Our job then in defending meaning will be to find the impetus and manifestations of these linguistic expectations, which suggests a handle on meaning.

Similarity

In truth, what makes meaning possible is my observation of your external state, and my reflection on my own internal states. Since we are quite similar beings, I can often ascertain someone's internal state by recognizing familiar patterns. I know that when I act in such a way it is because I am comfortable, and in such a way because I am afraid. This gives me a window into someone else's mind, one that is independent of language, if more rudimentary than it. It is through reflection on my own state, and the similarity between individuals in a linguistic community, which meaning can be based upon.

Even in this defense itself we see the poverty of meaning and the strength of skepticism. For there is no simple fact which makes clear that we are in fact similar beings. It also definitely the case that such a fact ought not to be linguistic in nature. Thus, even as I explain how I come to know meaning, and how I see myself reflected in my fellow speaker, I fall prey to the possibility of skepticism. This is because there is nothing laid out in all the facts about language which necessitates that there is in fact this perceived resonance.

One might rejoinder that our use of language itself indicates that we are similar beings. If we compare ourselves to other living things around us, none seem to speak as we do. Therefore, there is a single concrete fact establishing our similarity, and it is a language use fact to boot! Namely, the fact itself that we speak indicates that we share enough to have deep insight into one another's states. This response quickly crumbles, however, as it is clearly possible for language using beings to share little anatomy in common. For instance, if one species were land dwelling and the other aquatic, or one species surface dwellers and the other lived in deep caverns. Thus, there is little in the fact of language use itself to indicate similarity, nor anything else that can't again be viewed skeptically. The underlying issue here is that we don't have direct access to other minds.

And so again we are at an impasse. The naïve view of meaning refers to our commonality as a source of information about meaning, but this isn't encoded as a single simple fact, and so it is difficult to use in an argument. Even though I do indeed know my fellow language speaker through introspection and pattern recognition, I cannot prove this fact to a skeptic. Thus a skeptic may cast doubt on whether I in fact have any insight at all into other minds, and ask of me how I know that I ever make sense at all to others. Maybe it is the case that their internal states are radically different than mine, and I simply am projecting my subjective meanings onto them, when the use facts alone do not determine any such meaning assignments. Even though I have reason to believe in meaning, I do not have an argument for it.

The intangibility of similarity then accounts for the possibility of skepticism. The point is that the things which confirm meaning to me come in multitudes of tiny subtleties. If I am to point to any one of them, it will quickly turn to dust in my hand. Yet taken as a whole, they continually confirm my picture of other minds. I can point to little but the entirety of my experience as evidence. This leaves me in the odd position of being extremely convinced that my communication is meaningful, while being completely unable to make explicit what it is that makes me think so. At the end of the day I just sit with my conviction, resigned to allow skepticism to persist.

Complexity

One perhaps comforting thought is that language ought to allow such skeptical arguments. Since language is such a complicated process, the affirmation of meaning naturally doesn't lie in any of its individual smaller parts. Language is indeed a fickle thing, full of self-reference, and designed to be isolated from its content. The business of bridging the mind gap is indeed difficult. In this process of abstracting, there are no guarantees that things fall as they should. Even though at the end of the day I can restore meaning and explicate in great length how it is that I come to think of meaning as I do, there are so many twists and turns along the way, and so layered are the arguments, that I will never be free of the skeptical arguments. This is just the price we pay for as wonderful a tool of language, by necessity it is a dance of habit with thought, wrought so that one can never be sure what they are trying to say.