(1) Brandom begins by describing three methodological distinctions between possible approaches to philosophy of language.

   (a) Between what he calls the categorial *bottom-up* and *top-down*,

   (b) between what he calls the semantically *representational* and *expressive*, and

   (c) between a *empiricist* and *rationalist* choice of conceptual paradigm.

(2) The categorial dimension is concerned with what we should take as the first object of analysis which, then, provide the basis to explain the rest.

   (a) A bottom-up approach takes individual words to be the prime objects of analysis. We begin by considering individual words, terms and predicates, and analyse how they function completely. Afterwards, we determine how our theories of these individual, types of, words can be combined to explain full sentences and more complex linguistic phenomena. These are, essentially, approaches which take the principle of compositionality as a central feature of language, which provides us with a methodology for constructing semantic theories.

   (b) A top-down approach, on the other hand, begins with the analysis of complete sentences. It attempts to figure out the contents of entire sentences and only then tries to use that information to give an account of its constituent parts.

   We've seen paradigmatic examples of bottom-up approaches from all the theorists mentioned when discussing the theory of reference as well as from proponents of truth-conditional semantics. Top-down approaches, on the other hand, we've seen from both Grice and the later Wittgenstein.

(3) The semantic dimension is concerned with whether what's being represented or what's being expressed takes explanatory priority.

   (a) A representational approach begins with trying to explain what's being picked-out or referred to by expressions and then using that as part of the analysis of what the expressions mean. This is what he sometimes in the paper calls "of intentionality" as in 'What does the sentence speak of?'.

   (b) Conversely, an expressive approach begins with what is being expressed or said by an expression and then using that to explain what is being represented by it. This is what he sometimes in the paper calls "that intentionality" as in 'What does the sentence say that?'.

   Essentially, one can think of this dimension as being about whether the account takes meaning to be explained in terms of reference or reference in terms of meaning. Truth-conditional theories, for example, are normally representational approaches whereas Fregean semantics is expressive.

(4) The choice of conceptual paradigm has to do with what one thinks is the primary kind of explanation for the content of linguistic expressions.

(a) An empiricist approach to concepts is one where observations are the key to defining their content. That is, what is central to the concept of 'Red' is that we apply it in response to red things.

(b) A rationalist approach to concepts is one where the relationship between different concepts is key to defining their content. So, what is central to the concept 'Red' is that if we know something is red then we can conclude that it is coloured and that it is not green.

So, the verificationists are a typical example of philosophers who've taken an empiricist approach to concepts. The rationalist approach is, as Brandom notes, more associated with how we define logical symbols like $\land, \lor$, and $\to$ in terms of the inferential-rules they follow. As a logical connective, 'and' ($\land$) has its meaning fully specified by the following introduction and elimination rules.

$$\frac{P \quad Q}{P \land Q}\,\text{I} \qquad \frac{P \land Q}{P}\,\text{E} \qquad \frac{P \land Q}{Q}\,\text{E}$$

(5) So, Brandom categorises his approach as top-down, expressivist, and rationalist. That is, he takes full sentences as the first object of analysis, their meaning as explanatorily primary over their reference, and that what is key to identifying the conceptual content of sentences is how they relate to other sentences in reasoning.

(6) He also accepts that although he takes these features as being the first part of explaining linguistic phenomena, his theory still carries the burden of being able to explain how individual words work, reference, and observational uses of concepts in terms of his choice of primary explainers.

(7) The rest of Section II is incredibly densely written even to professional philosophers of language. It only contains an attempt at motivating inferentialism through its (at this point purported) ability to explain what we mean by modal, probabilistic, and normative words in terms which do not rely on unexplained metaphysical concepts (such as possible worlds).

(8) In section III, Brandom sets out to distinguish three types of inferentialism. The first point is to be clear about what he means by an *inferential connection* between sentences. The idea is that there is an inferential connection between the sentences

$$\frac{\text{This apple is red.}}{\text{This apple is coloured.}}$$

in the sense that it would be *correct* to conclude to conclude the latter from the former. So, it's not a matter of whether we actually do or don't make this inference, but about whether it would or wouldn't be correct to do so.

(9) He also wants to distinguish between a *narrow* and *wide* construal of inferences. On a narrow construal, the only inferences which are correct are those which are so in virtue of their *logical form* and the only thing which can serve as the premise or conclusion of an inference is a sentence. But on a *wide construal* of inference, we also accept the following:

(a) Inferences can be correct in virtue of their conceptual content and not just their logical form. That is, on a wide construal

$$\frac{A \text{ is to the west of } B}{B \text{ is to the east of } A}$$

is a valid inference, even without an additional premise which explicitly says that 'If A is to the west of B, then B is to the east of A.' (which would be required to make the argument *logically* valid). Brandom calls these kinds of inferences based on content *material inferences*.

(b) On a wide construal of inferences we also accept that there are some correct inferences based on material *incompatibilities* between claims, such as

$$\frac{\text{That's square.}}{\text{That's not triangular.}}$$

(c) Inferences whose premises are circumstances of observation as well as inferences whose conclusions are the performance of actions. Writing within brackets to express a non-linguistic circumstance or action, this would be inferences like:

$$\frac{\text{[Standing in front of a red apple.]}}{\text{That is red.}} \qquad \frac{\text{I'm a bank employee}}{\text{[I put on a necktie.]}}$$

(10) The three kinds of inferentialism, then, are as follows:

(a) *Weak inferentialism* is the thesis that inferential connections between sentences are necessary for them to have the content they do.

(b) *Strong inferentialism* is the thesis that inferential connections broadly construed are both necessary and sufficient to determine the content of sentences.

(c) *Hyperinferentialism* is the thesis that inferential connections narrowly construed are both necessary and sufficient to determine the content of sentences.

As Brandom says, almost everyone would endorse weak inferentialism (and given, say, a truth-theoretic view of meaning, it is satisfied since if W, W' are the truth-conditions — sets of possible worlds — for P, Q, then Q follows from P if and only if every world in W' is in W.) Hyperinferentialism, he thinks, is true for logical vocabulary, but not language in general. What he endorses is strong inferentialism.

(11) The rest of section III is Brandom sketching how this kind of approach can be go on to meet the explanatory challenges which I mentioned that he acknowledges in (6). His description here is incredibly sparse (he uses three paragraphs to present 172 pages of material) and I don't expect you to follow it. I barely do, and I've read the damn book and know what he's referring to.

(12) In section IV he gets to one of the main challenges offered against this kind of approach, namely the question of which of the many inferences that a sentence can figure in make up parts of its meaning and which do not? For instance, is the following inference correct because of the meaning of the two sentences?

$$\frac{\text{Lightning was seen now.}}{\text{Thunder will be heard soon.}}$$

(13) This issue is, as Brandom says, essentially parallel to Quine's challenge to the distinction between analytic and synthetic truths. The inferences which make up the meaning of a sentence can be considered analytic inferences, whereas those which just happen to be correct can be considered synthetic. To fully learn how to use a sentence, then, would be to learn all the analytic inferences it figures in.

(14) Brandom discusses two possible answers to this question, before settling on the latter.

    (a) Sellars offers a solution in terms of saying that analytic inferences are precisely those which are *counterfactually robust*. That is, those inferences which aren't just correct in the actual world, but are also correct in every possible world. For example, the inference expressed by

<div align="center">If a beer is in my fridge, then it is German.</div>

is correct, but it isn't counterfactually robust because it's it's not correct that

<div align="center">If a beer would be in my fridge, then it would be German.</div>

So, on Sellars' answer to the question, the inference from being in my fridge to being German is not meaning-constitutive for the concept 'Beer'. On the other hand,

<div align="center">If something is wine, then it is made from grapes.</div>

is counterfactually robust, because it's correct to say that

<div align="center">If something would be wine, then it would be made from grapes.</div>

So, this inference is part of the meaning of the concept 'Wine' according to Sellars.

    (b) But this is not the approach Brandom opts for. Instead, he adopts the view that there isn't an important distinction between analytic and synthetic inferences. So, he thinks that *every* inference that a sentence figures takes part in making it mean what it does.

(15) This, as he notes, has the consequence that we can't really identify the meaning of a sentence in isolation. Brandom must accept that the entirety of a language as relevant to determining what a single statement means, just as Quine concludes from his rejection of an analytic-synthetic distinction for statements. This, as we've discussed earlier, raises the worry about whether a theory which has this *holistic* consequence about meaning could ever explain how we can communicate. Responding to this worry is the main thrust of section V.

(16) The first thing he discusses here is how Quine approached his version of the problem, which is by saying that communication and language learning proceed through shared reference. So, while two people might mean very different things when saying that

<div align="center">The sun is a visible thing.</div>

because they believe very different things about the sun, they are still referring to the same thing when saying it. And this, Quine thinks, is enough to explain communication and language learning. Note that this doesn't contradict Quine's view on the indeterminacy of reference and translation. His claim there is that when faced with a speaker of a different language we can't determine how the reference of expressions in that persons language compare to the reference of expressions in ours. He does not claim that two individuals couldn't get a language started by jointly stipulating that certain sounds correspond to certain, shared, observations.

(17) But this approach is not open to Brandom, because it takes the representational rather than expressive as the explanatorily prior semantic phenomena. And, as I said in (5), Brandom wants to take the expressivist approach. So, while he does claim that his account explains how we do share reference (note: what he says on pp. 664–665 is not really an explanation, it just points to the one he gives in the massive book), he wants to explain how this follows from the inferential account of meaning and, hence, it would be circular for him to use it to support that account.

(18) Next, he discusses an option which he doesn't ultimately go for, where we would explain communication by partially attaching the same meaning to an expression. For an inferentialist, this would be when we agree on the correctness of some, but not all, of the inferential relations the expression figures in.

(19) The second option, which he does endorse, rejects what he calls the Lockean view of communication. This is the idea, which we've talked about previously, that communication functions by transmitting a specific meaning from speaker to hearer. This is supposed to proceed via encoding that meaning in the uttered expression which the hearer can then decode in virtue of knowing the correspondence relation between expressions and meanings. It is this picture, Brandom claims, which makes holism seem like a problem for communication, since we would have to be in complete agreement about the meaning of *every* sentence in order to be in agreement about the meaning of *any* sentence.

(20) But we can drop this view of communication. Instead, we can think of communication as a practical ability to co-ordinate beliefs and actions across different perspectives on what our words and sentences mean. On this kind of view, communication is like trying to dance with someone even though you might agree on what the correct steps are. But because your respective conceptions of the steps are close enough, and you're attentive to the other persons moves, we can manage to successfully perform a dance anyway. And the idea is that we can provide a similar account of communication, where we're responsive to differences in our usages and manage to successfully co-operate about our beliefs and actions, then we've managed to communicate even if we can't ever transmit belief.

(21) Brandom, however, takes a third option to be most significant. He makes yet another distinction based on whether one thinks of concepts in a *Cartesian* or *Kantian* way.

  (a) Concepts in a Cartesian sense are a kind of mental particulars. They are some sort of structure or thing in the mind of the concept-haver. So, understood like this, you and I can have different concepts of 'Red' because your mental particular stands in different inferential relations than mine, in the sense that we judge different inferences to be correct. As such, the worry that holism prevents communication is a real problem, at least if we retain the Lockean view.

  (b) Concepts in a Kantian sense are instead to be understood as intersubjective norms or rules. That is, while you and I can have different *conceptions* of the rules governing the concept 'Red' — again, that means disagreeing on which inferences involving it are correct — there still is an intersubjectively shared rule which determines when we are right or not. Communication, Brandom thinks, is about navigating the game governed by these rules whether we are aware of them or not. In that sense, what I say when I use the concept 'Red' is not determined by *my* conception of the rules governing it, but rather by the actual rules; although if someone wants to understand what I think I've said, then my conception is the determining factor.

(22) Section VI turns to the related question of compositionality. The challenge for the inferentialist is to explain how we can produce and understand new and unfamiliar sentences. Normally, as Brandom says, this is explained by a principle of compositionality, but this approach is in tension with the top-down categorical approach Brandom has committed himself to.

(23) He begins by trying to resolve this tension by arguing that one can accept both that compositionality explains productivity and that the content of complete sentences is explanatorily prior to that of individual components, like terms and predicates. The idea, which he attributes to Dummett and Frege, is to say that we can analyse the meaning of, say, a term through it's contribution to the sentences it figures in. For example, we can say that two singular terms s, t have the same *indirect inferential role* if every correct inference involving a sentence containing s would still be correct when we replace s with t. As such, we can construct something that looks a lot like the meaning of a singular term which respects compositional structure, even if we start from the meaning of full sentences.

(24) The details of this proposal are a lot more complicated, and Brandom doesn't really explain them here, so, the main point is to show that inferentialism is compatible with a principle of compositionality, even if it has to be explained in a less direct way. The rest of the section sketches some of the ways that the proposal is supposed to go through but, again, it is trying to compress to many details in too little space. I won't discuss it more here and I don't expect you to get much out of reading it.

(25) The reason we read this paper in the course, then, is to get an inside look into one of the more impressive contemporary use-theoretic approaches to meaning and how they can resist the criticism that are often levelled at Wittgenstein. Brandom, I think, provides a useful discussion of just how different this kind of approach can be from more traditional propositional or truth-conditional semantics. In the conclusion, he also makes the important point that to analyse theories which have a radically different structure it's important not to consider them piecemeal. As we can see from the myriad of distinctions he's made, a lot of the moves an inferentialist make in response to criticism look like they fail unless one also let's go of some unspoken commitment of the traditional view.