

Excerpts from
*PHILOSOPHICAL
 INVESTIGATIONS*

Ludwig Wittgenstein

Source: *Philosophical Investigations* (Oxford, 1953): selected, edited and with commentary by Yorick Wilks, Oxford, 2005.

*When reading Wittgenstein—or at least his mature work—a number of unmentioned presences have to be kept in mind at all times. The principal one is Wittgenstein's early self, and his 'picture theory of truth'. Much of the motivation of the *Philosophical Investigations* (Blackwell, 1953) was to set out why that view and its associated doctrines were wrong. Also in the background are Tarski and Carnap, who still advocated formalist views long after Wittgenstein had given them up. Most of the views attacked in *Philosophical Investigations* were held in one form or another by Tarski. The change in Wittgenstein was that he ceased to believe that words in natural language had meaning chiefly because they pointed at objects in the world and that sentences were true because they matched up to the world in some direct one-to-one way. He became more and more convinced that what was important about language was its deep grammatical forms, and it was from here that the metaphor of 'depth' in modern linguistics took off. Wittgenstein always resisted any attempt to formalize a theory of these deep forms, and there is no point in imagining that he would have been deliriously happy had he lived to see modern linguistics and AI as alternatives to the logical paradigm. But many of the concerns of modern linguistics, NLP, and AI are already there in his work, and his line of thinking is a powerful antidote to the naive errors with which the subject is still riddled.*

*Wittgenstein had a peculiar style: his work, early and late, takes the form of a series of numbered remarks. Some of these were arranged in their present order after his death, in the early fifties, by editors. The remarks are not themselves arranged neatly under headings, and reading Wittgenstein therefore takes the form of tracing connexions through the remarks for oneself. The quotations that follow are from his *Philosophical Investigations*, with the*

exception of a few from the *Philosophical Remarks (PR)*. Each section that follows consists of a thesis, then one or more quotations, followed by comment.

Here is an epigraphic quote for all that follows:

§ 122. A main source of our failure to understand is that we do not command a clear view of the use of our words—our grammar is lacking in this sort of perspicuity. A perspicuous representation produces just that understanding which consists in ‘seeing connexions’. Hence the importance of finding and inventing intermediate cases. The concept of a perspicuous representation is of fundamental significance for us. It earmarks the form of account we give, the way we look at things.

i. Reference

Thesis: words do not in general have meaning in virtue of pointing at objects in the real world (or at ‘conceptual objects’ either).

§ 2. That philosophical concept of meaning [i.e. of meaning as thing] has its place in a primitive idea of the way language functions, But one can also say that it is the idea of a language more primitive than ours.

§ 13. When we say: “Every word in language signifies something” we have so far said nothing whatever; unless we have explained exactly what distinction we wish to make. (It might be, of course, that we wanted to distinguish the words of [some] language from words ‘without meaning’ such as occur in Lewis Carroll’s poems, or words like ‘Lilliburlero’ in songs.

§ 30. So one might say: the ostensive definition [i.e. pointing] explains the use—the meaning—of the word when the overall role of the word in language is clear. Thus if I know that someone means to explain a colour-word to me the ostensive definition “That is called ‘sepia’” will help me to understand the word.

§ 35. There are, of course, what can be called “characteristic experiences” of pointing to (e.g.) the shape. For example, following the outline with one’s finger or with one’s eyes as one points.—But this does not happen in all cases in which I ‘mean the shape’ and no more does any other one characteristic process occur in all these cases. Besides, even if something of the sort did recur in all cases, it would still depend on the circumstances—that is, on what happened before and after the pointing—whether we should say “he pointed to the shape and not to the colour”.

Comment: Wittgenstein is arguing that pointing or referring is in principle a vague activity. It can only be made clear by explaining from within the

language what we are pointing at—i.e. useful pointing already assumes the whole language. Hence it is not that pointing explains how we mean, as the formalists thought when they defined the denotations of their symbols as objects, or sets of objects, because, argues Wittgenstein, the pointing presumes upon the language rather than explains it.

Wittgenstein says we could have a language based on the referential notion (§2), but it would be a language more primitive than what we call natural language. The general point here is very close to the one made later by Quine (*Word and Object*, pp. 29–46) when arguing for the essential ambiguity of terms like ‘Gavagai’ if used ostensively by people with whom we could not communicate because we shared no language, what we can think of as the ‘anthropologist scenario’. The problem with this version of Wittgenstein’s point is that human experience shows that total strangers can learn languages from such unpromising starts, so his point must apply to individual acts in such situations.

ii. Mini languages and language games

Thesis: we can construct mini-languages obeying any rules we like, and we can think of them as games. The important question is whether these games are sufficiently like the whole game of natural language. This question does not have a definite answer any more than this question: “Can one play chess without the queen?” Wittgenstein attributes the ostensive or ‘pointing’ view of meaning (in (i) above) to St. Augustine and then proceeds to construct a mini-language of commands and objects like ‘block’, ‘slab’ and colours like ‘red’, etc.

§ 2. Let us imagine a language for which the description given by Augustine is right. The language is meant to serve for communication between a builder A and an assistant B. A is building with building stones: there are blocks, pillars, slabs and beams. B has to pass the stones, and that in the order in which A needs them. For this purpose they use a language consisting of the words ‘block’, ‘pillar’, ‘slab’, ‘beam’. A calls them out; B brings the stone which he has learnt to bring at such-and-such a call. Conceive this as a complete primitive language.

§ 3. Augustine, we might say, does describe a system of communication; only not everything that we call language is this system. And one has to say this in many cases where the question arises, “Is this an appropriate description or not?” The answer is: “Yes, it is appropriate, but only for this narrowly circumscribed region, not for the whole of what you were claiming to describe”. It is as if someone were to say: “A game consists in moving objects about on a surface according to certain rules . . .”—and we replied: You seem to be thinking of board games, but there are others. You can make your definition correct by expressly restricting it to those games.

Comment: *This mini-language that Wittgenstein constructs with 'block', 'slab' and commands may remind readers of Winograd's mini-language (Dreyfus, 1972) inherited from MIT 'table-top robotics': it had a box, a block, a sphere and so on. The parallel is a fair one in many ways, and Wittgenstein can be seen as presenting the dangers of taking a mini-language with certain properties (definite reference to objects, for example) and assuming that they are properties of the whole of natural language. One could say that it is not clear how, or whether, a Winogradian system could function in a world without definite locatable objects, such as the world of newspaper articles or of this book.*

It should not be thought that Wittgenstein is a defender of linguistic primitives, or primitives of any sort. Indeed, one of the attractions to him of his 'truth-table' method of presenting the Propositional Calculus was that it avoided the more conventional form in terms of primitive formulas, like $P \text{ IMPLIES } (Q \text{ OR NOT-}Q)$ from which all other true formulas could be derived. Yet, one could argue that a large part of what Wittgenstein found objectionable about the notion of 'primitive' in logic was the idea that there is a right set of them, if only we could discover it, and which provides an infallible starting point. But in the case of semantic primitives, it is still possible (Wilks, 1986) to use them without claiming that there is a single right set of them, as Schank did.

iii. Family resemblances and boundaries

Thesis: *the conventional notion of a concept is wrong: namely the view that a concept relates in some way to the qualities or characteristics that all things falling under the concept have, as when one might claim in a simple-minded way that everything that is an arch has such-and-such properties. Wittgenstein takes the concept of a game and argues that one could not define a game by necessary and sufficient qualities: for any proposed necessary characteristic of being a game, Wittgenstein claims he can think of a game that does not have the characteristic. Patience (solitaire), for example, is not competitive, and so on. Along these lines it has been argued that entities under a concept form a something more like a family, for some members of a family share the characteristic long nose, while some share the characteristic tiny feet, but there need be no characteristic they all share for them to have a 'family resemblance'. The moral is that there are no firm boundaries to concepts, nor are there to linguistic usage, nor to the application of linguistic rules.*

§ 66. Consider for example the proceedings that we call "games". I mean board-games, card-games, ball-games, Olympic games, and so on. What is common to them all?—Don't say: "There must be something common, or they would not be called games"—but look and see whether there is anything common to all. For if you look at them you will not see something that is common to all, but similarities, relationships, and a whole series of

them at that. To repeat: don't think, but look!—Look for example at board-games, with their multifarious relationships. Now pass to card-games; here you find many correspondences with the first group, but many common features drop out, and others appear. When we pass next to ball-games, much that is common is retained, but much is lost.—Are they all 'amusing'? Compare chess with noughts and crosses. Or is there always winning and losing, or competition between players? Think of patience. In ball games there is winning and losing; but when a child throws his ball at the wall and catches it again, this feature has disappeared. Look at the parts played by skill and luck; and at the difference between skill in chess and skill in tennis. Think now of games like ring-a-ring-a-roses; here is the element of amusement, but how many other characteristic features have disappeared! And we can go through the many, many other groups of games in the same way; can see how similarities crop up and disappear.

And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail.

§ 67. I can think of no better expression to characterize these similarities than "family resemblances"; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc., etc. overlap and criss-cross in the same way.—And I shall say: 'games' form a family.

And for instance the kinds of number form a family in the same way. Why do we call something a "number"? Well, perhaps because it has a—direct—relationship with several things that have hitherto been called number; and this can be said to give it an indirect relationship to other things we call the same name. And we extend our concept of number as in spinning a thread we twist fibre on fibre. And the strength of the thread does not reside in the fact that some one fibre runs through its whole length, but in the overlapping of many fibres.

But if someone wished to say: "There is something common to all these constructions—namely the disjunction of all their common properties"—I should reply: Now you are only playing with words. One might as well say: "Something runs through the whole thread—namely the continuous overlapping of those fibres".

§ 68. "All right: the concept of number is defined for you as the logical sum of these individual interrelated concepts: cardinal numbers, rational numbers, real numbers, etc.; and in the same way the concept of a game as the logical sum of a corresponding set of sub-concepts."—It need not be so. For I *can* give the concept 'number' rigid limits in this way, that is, use the word "number" for a rigidly limited concept, but I can also use it so that the extension of the concept is *not* closed by a frontier. And this is how we

do use the word “game”. For how is the concept of a game bounded? What still counts as a game and what no longer does? Can you give the boundary? No. You can *draw* one; for none has so far been drawn. (But that never troubled you before when you used the word “game”.)

“But then the use of the word is unregulated, the ‘game’ we play with it is unregulated.”—It is not everywhere circumscribed by rules; but no more are there any rules for how high one throws the ball in tennis, or how hard; yet tennis is a game for all that and has rules too.

§ 69. How should we explain to someone what a game is? I imagine that we should describe games to him, and we might add: “This and similar things are called ‘games’”. And do we know any more about it ourselves? Is it only other people whom we cannot tell exactly what a game is?—But this is not ignorance. We do not know the boundaries because none have been drawn. To repeat, we can draw a boundary—for a special purpose. Does it take that to make the concept usable? Not at all. (Except for that special purpose.) No more than it took the definition 1 pace = 75 cm. to make the measure of length ‘one pace’ usable. And if you want to say “But still, before that it wasn’t an exact measure”, then I reply: very well, it was an inexact one. Though you still owe me a definition of exactness.

§ 70. “But if the concept ‘game’ is uncircumscribed like that, you don’t really know what you mean by a ‘game’.” When I give the description: “The ground was quite covered with plants”, do you want to say I don’t know what I am talking about until I can give a definition of a plant?

§ 71. One might say that the concept game is a concept with blurred edges. “But is a blurred concept a concept at all?” Is an indistinct photograph a picture of a person at all? Is it even always an advantage to replace an indistinct picture by a sharp one? Isn’t the indistinct one often exactly what we need?

§ 76. If someone were to draw a sharp boundary I could not acknowledge it as the one that I too always wanted to draw, or had drawn in my mind. For I did not want to draw one at all. His concept can then be said to be not the same as mine, but akin to it. The kinship is that of two pictures, one of which consists of colour patches with vague contours, and the other of patches similarly shaped and distributed, but with clear contours. The kinship is just as undeniable as the difference.

§ 84. I said that the application of a word is not everywhere bounded by rules. But what does a game look like that is everywhere bounded by rules? whose rules never let a doubt creep in, but stop up all the cracks where it might?—Can’t we imagine a rule determining the application of a rule, and a doubt which it removes—and so on?

§ 80. I say “There is a chair.” What if I go up to it, meaning to fetch it, and it suddenly disappears from sight? “So it wasn’t a chair, but some kind of illusion.” But in a few moments we see it again and are able to touch it and so on.—“So the chair was there after all and its disappearance was some kind of illusion.” But suppose that after a time it disappears again, or seems to disappear. What are we to say now? Have you rules ready for such cases—rules saying whether one may use the word ‘chair’ to include this kind of thing? But do we miss them when we use the word ‘chair’; and are we to say that we do not really attach any meaning to this word, because we are not equipped with rules for every possible application of it?

§ 88. If I tell someone “Stand roughly here,” may not this explanation work perfectly? And cannot every other one fail too? But isn’t it an inexact explanation? Yes; why shouldn’t we call it “inexact”? Only let US understand what “inexact” means. For it does not mean “unusable”.

§ 99. The sense of a sentence—one would like to say—may, of course, leave this or that open, but the sentence must nevertheless have a definite sense. An indefinite sense—that would really not be a sense at all. This is like: An indefinite boundary is not really a boundary at all. Here one thinks perhaps: if I say “I have locked the man up fast in the room—there is only one door left open”, then I simply haven’t locked him in at all; his being locked in is a sham. One would be inclined to say here: “You haven’t done anything at all”. An enclosure with a hole in it is as good as none—But is that true?

§ 100. “But still, it isn’t a game, if there is some vagueness in the rules.”—But does this prevent its being a game?—“Perhaps you’ll call it a game, but at any rate it certainly isn’t a perfect game.” This means: it has impurities, and what I am interested in at present is the pure article. —But I want to say: we misunderstand the role of the ideal in our language. That is to say: we too should call it a game, only we are dazzled by the ideal and therefore fail to see the actual use of the word ‘game’ clearly.

§ 133. It is not our aim to refine or complete the system of rules for the use of our words in unheard-of ways.

§ 249. Lying is a language-game that needs to be learned like any other one.

Comment: There are many large issues touched on in these quotations, and the other paragraphs that surround them in the original. The central theme here is that natural language has an essential vagueness in its everyday employment, one that does no harm, and which in no way interferes with our understanding each other. It is, however, a feature that no one at the moment has any idea how to put into a computational system. On the contrary, virtually all NLP

systems assume that there is a correct and contentful interpretation to any piece of language use, and that we can decide by fairly simple rules what that interpretation is. The considerations Wittgenstein brings forward in these quotations suggest that in real life it is not so simple.

It is certainly not true, as is often assumed by many NLP and AI workers, that language is fundamentally a mechanism by which experts in various fields convey information to each other. If that assumption was true then 99% of all actual human communication would be in some way inadequate—and that just cannot be so. The quotations suggest that a feature a satisfactory language system would have to embody is the ability to redraw its own boundaries, and to be as vague or specific as necessary, in the way that we can.

To widen our scope slightly, these considerations might suggest that a theory of language can never be, in a straightforward way, a scientific theory†: that is to say, a body of rules that decides, rightly or wrongly, of any utterance, what its correct structure is. But that is the assumption on which most NLP is based! A fuller discussion of §133 might lead us to question the whole possibility and usefulness of ‘ingenious knockdown counter-examples’ to theories of language, which linguistics has made a speciality, but which has no place in more recent statistical language processing.

iv. The linguistic whole and confronting the world

Thesis: a language is a whole and does not confront the world sentence by sentence for the testing of the truth or falsity of each individual part.

§ 199. To understand a sentence means to understand a language.

Comment: This thesis is clearly incompatible both with Wittgenstein’s own early ‘picture theory of truth’ and with any theory like Montague’s, where the assumption is precisely that each sentence of a language has its truth (and its meaning) tested individually and in isolation. One could argue that Wittgenstein’s view is not at all inconsistent with a standard view of scientific truth, where sentences such as ‘This particle has spin $1/2$ ’ or ‘Rats are carriers of plague’ are not tested directly but belong only within large systems of inference that must be tested indirectly if at all. That is to say that sentences like those two can only be understood within a wider theory which, in its turn, explains complex notions like **spin**, **particle** and **carried by**. This point is also close to the heart of Quine’s philosophy, and for him followed from his analysis of the Analytic-Synthetic dichotomy, such that if sentences could not simply be assigned to one of these classes then how much a sentence had of one property or the other would depend on its position and role in a scientific theory.

The notion that we can only understand on the basis of a whole language is clearly more attractive than the distant alternative of understanding within ‘block and slab’ micro-worlds. But it, too, has its dangers: if taken far enough,

it can lead to the view that there can be no significant generalizations about language at all, because each use of each sentence has a special relation to the language as a whole. And Wittgenstein has sometimes been accused of holding this view.

What seems more likely to have been his position was an intermediate one; namely, that there are islands of discourse each with its own criteria of inference, truth and so on, analt is with respect to these (rather than to a micro-world or to the whole language) that utterances are to be understood. The notion of an ‘island of discourse’ is not a self-explanatory one but, roughly speaking, it means an area defined by subject matter (say, history or quantum physics), but still wide enough to have all the features of a full natural language, in a way that a ‘slab and block’ micro-world does not, nor does a narrow domain application of the kind that constitutes much of applied NLP, e.g. airline reservation worlds.

v. Logicians have a false picture of how language is

Thesis: logicians think that language is like their favorite calculus, but they are quite wrong. Moreover, it is language itself and its use that is the standard for testing disputes that arise, not what logicians dictate.

§ 81. F. P. Ramsey once emphasized in conversation with me that logic was a “normative science”. I do not know exactly what he had in mind, but it was doubtless closely related to what only dawned on me later: namely, that in philosophy we often compare the use of words with games and calculi which have fixed rules, but cannot say that someone who is using language must be playing such a game. But if you say that our languages only approximate to such calculi you are standing on the very brink of a misunderstanding. For then it may look as if what we are talking about is an ideal language. As if our logic were, so to speak, a logic for a vacuum. Whereas logic does not treat of language—or of thought—in the sense in which a natural science treats a natural phenomenon, and the most that can be said is that we construct ideal languages. But here the word ‘ideal’ is liable to mislead, for it sounds as if these languages were better, more perfect, than our everyday language, and as if it took the logician to show people at last what a proper sentence looked like.

All this, however, can only appear in the right light when one has attained greater clarity about the concepts of understanding, meaning, and thinking. For it will then also become clear what can lead us (and did lead me) to think that if anyone utters a sentence and means or understands it he is operating a calculus according to definite rules.

§ 91. But now it may come to look as if there were something like a final analysis of our forms of language, and so a single completely resolved form

of every expression. That is, as if our usual forms of expression were, essentially, unanalyzed, as if there were something hidden in them that had to be brought to light. When this is done the expression is completely clarified and our problem is solved. It can also be put like this: we eliminate misunderstandings by making our expressions more exact, but now it may look as if we were moving towards a particular state, a state of complete exactness, and as if this were the real goal of our investigation.

§ 101. We want to say that there can't be any vagueness in logic. The idea now absorbs us, that the ideal 'must' be found in reality. Meanwhile we do not as yet see how it occurs there, nor do we understand the nature of this 'must'. We think it must be in reality, for we think we already see it there.

§ 115. A picture held us captive. And we could not get outside it, for it lay in our language and language seemed to repeat it to us inexorably.

Comment: This thesis clearly clashes head on, not only with Montague, but also with those logicians who believe in the useful applicability of the Predicate Calculus to language. Notice that Wittgenstein is not claiming that the logicians are being inconsistent, as between their beliefs and the way they talk every day of their lives, any more than phlogiston theorists in the Eighteenth century were being inconsistent when they speculated on the phlogistic process (the reverse of chemical oxidation) while their lungs actually kept them alive by oxidation processes. They were simply describing phenomena they had not examined.

vi. Understanding is not a feeling

Thesis: we have the idea that understanding something involves, or is associated with, a special feeling of being right. But the tests of our being right are quite different from the feeling.

PI, page 59, footnote. (a) *Understanding a word:* a state. But a mental state?—Depression, excitement, pain, are called mental states. Carry out a grammatical investigation as follows: we say:

"He was depressed the whole day."

"He was in great excitement the whole day."

"He has been in continuous pain since yesterday."

We also say: "Since yesterday I have understood this word." *Continuously*, though?—To be sure, one can speak of an interruption of understanding. But in what cases? Compare: "When did your pains get less?" and "When did you stop understanding that word?"

§ 122. A perspicuous representation produces just that understanding which consists in 'seeing connexions'.

§ 139. When someone says the word 'cube' to me, for example, I know what it means. But can the whole use of the word come before my mind, when I understand it in this way?

Well, but on the other hand isn't the meaning of the word also determined by this use? And can these ways of determining meaning conflict? Can what we grasp in a flash accord with a use, fit or fail to fit it? And how can what is present to us in an instant, what comes before our mind in an instant, fit a use?

What really comes before our mind when we understand a word?—Isn't it something like a picture? Can't it *be* a picture?

Well, suppose that a picture does come before your mind when you hear the word 'cube', say the drawing of a cube. In what sense can this picture fit or fail to fit a use of the word 'cube'?—Perhaps you say: "It's quite simple; if that picture occurs to me and I point to a triangular prism for instance, and say it is a cube, then this use of the word doesn't fit the picture."—But doesn't it fit? I have purposely so chosen the example that it is quite easy to imagine a method of projection according to which the picture does fit after all.

The picture of the cube did indeed suggest a certain use to us, but it was possible for me to use it differently.

§ 151. But there is also this use of the word 'to know': we say "Now I know"—and similarly "Now I can do it:" and "Now I understand!"

Let us imagine the following example: A writes series of numbers down: B watches him and tries to find a law for the sequence of numbers. If he succeeds he exclaims: "Now I can go on!"—So this capacity, this understanding, is something that makes its appearance in a moment. So let us try and see what it is that makes its appearance here.—A has written down the numbers 1, 5, 11, 19, 29; at this point B says he knows how to go on. What happened here? Various things may have happened: for example, while A was slowly putting one number after another, B was occupied with trying various algebraic formulae on the numbers which had been written down. After A had written the number 19 B tried the formula $a_n = n^2 + n - 1$; and the next number confirmed his hypothesis.

§ 155. Thus what I wanted to say was: when he suddenly knew how to go on, when he understood the principle, then possibly he had a special experience—and if he is asked: "What was it? What took place when you suddenly grasped the principle?" perhaps he will describe it much as we described it above—but for us it is the circumstances under which he had such an experience that justify him in saying in such a case that he understands, that he knows how to go on.

Part II, p. 181. Even if someone had a particular capacity only when, and only: as long as he had a particular feeling, the feeling would not be the capacity.

The meaning of a word is not the experience one has in hearing or saying it, and the sense of a sentence is not a complex of such experiences.—(How do the meanings of the individual words make up the sense of the sentence “I still haven’t seen him yet”?) The sentence is composed of the words, and that is enough.

Part II, pp. 187–188. Are the words “I am afraid” a description of a state of mind?

I say “I am afraid”; someone else asks me: “What was that? A cry of fear; or do you want to tell me how you feel; or is it a reflection on your present state?”—Could I always give him a clear answer? Could I never give him one?

We can imagine all sorts of things here, for example:

“No, no! I am afraid!”

“I am afraid. I am sorry to have to confess it.”

“I am still a bit afraid, but no longer as much as before.”

“At bottom I am still afraid, though I won’t confess it to myself.”

“I torment myself with all sorts of fears.”

“Now, just when I should be fearless, I am afraid!”

To each of these sentences a special tone of voice is appropriate, and a different context. I would be possible to imagine people who as it were thought much more definitely than we, and used different words where we use only one.

We ask “What does ‘I am frightened’ really mean, what am I referring to when I say it?” And of course we find no answer, or one that is inadequate.

The question is: “In what sort of context does it occur?”

Comment: First Wittgenstein is making the point that it is dangerous to assess understanding other than in terms of actual and possible performances. Secondly, there is a very general theme running through Wittgenstein’s work about the relation of knowledge and understanding to performance and to what he calls the ability to go on, to continue (§151). This theme could be held to support those who go beyond the assertion that our understanding of words depends on our ability to use them and to perform with them, to the much stronger and less plausible claim that our understanding of language about physical processes (say, tying our shoelaces, or stacking blocks) is closely connected with (and may even require) our ability to carry out the corresponding task. That would mean that a computer could not understand language about, say, dining in a restaurant, unless it could itself dine in a restaurant. This is a complex issue, usually discussed under terms like ‘grounding’ or ‘situatedness’ and one where Wittgenstein’s explorations are essential background.

Thirdly, there is often confusion between (1) what the processes actually are in our heads in carrying out a task. (2) how we feel about what they are, and (3) what a computer should do in program terms to carry out the same task. These are three quite separate things and arguments connecting them can be dangerous, as Wittgenstein warns in §139.

Dreyfus (1972) has argued that AI is impossible because, to be intelligent, a computer would have to be exactly like us: bodies, feelings and all. He has often quoted Wittgenstein in support of this position, but it can equally well be argued that Wittgenstein’s clear distinction between understanding-as-performance (which AI workers believe a machine can have) and understanding-as-feeling (which no doubt only humans have) supports exactly the opposite position.

vii. Application justifies our structures

Thesis: the significance of a representational structure cannot be divorced from the process of its application to actual language.

§ 43–5. We cannot compare a picture with reality if we cannot lay it against reality as a measuring rod.

PR., p. 308. The rod must be in the same space as the object to be measured.

Comment: These two tiny quotes only hint at Wittgenstein’s discussion of an issue raised in (iv) above: how do our representational structures actually function if they cannot be ‘laid directly against the world they represent’. They cannot be laid on directly because they are not ‘in the same space’ that is to say the representation is in symbols and the world is the chairs, tables, ocean, love, guilt, and government buildings out there. Those who advocate Montague-type systems do seem to believe that in some way their representations are ‘laid against the world’ directly, via such constructs as his L-semantic notions of *set* and *function*.

viii. Real world knowledge and forms of life

Thesis: language understanding is not independent of very general inductive truths about our human experience.

§ 142. It is only in normal cases that the use of a word is clearly prescribed; we know, are in no doubt, what to say in this or that case. The more abnormal the case, the more doubtful it becomes what we are to say. And if things were quite different from what they actually are—if there were for instance no characteristic expression of pain, of fear, of joy; if rule became exception and exception rule; or if both became phenomena of roughly equal frequency—this would make our normal language games lose their point. The procedure of putting a lump of cheese on a balance and fixing the price

by the turn of the scale would lose its point if it frequently happened for such lumps to suddenly grow or shrink for no obvious reason. This remark will become clearer when we discuss such things as the relation of expression to feeling, and similar topics.

II. xii. If the formation of concepts can be explained by facts of nature, should we not be interested, not in grammar, so much as in nature; which forms the basis of grammar? Our interest certainly includes the way concepts answer to very general facts of nature. (Such facts as usually do not strike us because of their generality.) But our interest does not revert to these possible causes of concept-formation: we're not doing natural science, nor even natural history, since we can indeed construct fictitious natural history for our purposes.

*Comment: NLP/AI researchers in the knowledge-based understanding tradition sometimes emphasize very general knowledge of the world, like that of human wishes and desires, while others emphasize the role in understanding of very particular knowledge, such as that of airline reservations or birthday parties. But it could be argued that if this knowledge of, say, human desires is essential to our understanding of our language (as Wittgenstein suggests in § 142) then that knowledge may in a sense be knowledge of our language and not knowledge of the physical world at all. That is to say that, on this view, 'Humans try to get what they want' is not really a descriptive fact about people, but is a disguised way of stating what we mean by **get** and **want**, in that we would not admit that someone really wanted something unless he tried to get it. One can see here in Wittgenstein the beginning of the rather simpleminded 'folk psychology' position of Dennett, which has also emphasized the continuity of human language and our understanding of ourselves.*

In conclusion, it must be admitted that Wittgenstein is a frustrating author: at times he writes so that he can be taken both ways on important questions. The value of his view of language, as Max Black put it, is that it contains the means for its own supplementation.

The contrast between Wittgenstein and formal theorists of language such as Montague comes down to two issues: 1) Is there a hidden structure to natural language? and 2) Is natural language itself, and its use, to be the final court of appeal? Montague gives a yes to the first and proffers a simple logic as the hidden structure; as to the second, his choice of examples, far from the concerns of ordinary speakers, such as the readings of 'Every man loves some woman' suggests that his answer is no. Wittgenstein's answer to the first is complex: his use of deep grammatical structures suggests that he thought there was structure but not one to be revealed by simple techniques, like logic, whose interests were really always somewhere else. It is, after all, the different structures of 'Every man loves some woman' as they can participate in proofs in the Predicate Calculus that interest the logician. The ordinary speaker rarely, if ever, sees that there is a second interpretation of the sentence.

References

- Hubert Dreyfus, *What Computers Can't Do: A Critique of Artificial Reason*, New York: Harper & Row, 1972.
 Terry Winograd, *Understanding Natural Language*, Academic Press, 1972.
 Wilks, Y. (1986) Primitives, In S. Shapiro (ed.), *Encyclopedia of Artificial Intelligence*, New York: Wiley.