

On Self-Referential Pseudo-Paradoxes

Reply to Douglas Hofstadter

Copyright Flash qFiasco 2010
<http://www.flashq.org>

Douglas Hofstadter has written:

Mr Flash qFiasco insisted that a sentence cannot say what it shows. The former concerns only its content, which is supposedly independent of how it manifests itself in print, while the latter is a property exclusively of its form, that is, of the physical sentence only when it is in print. This distinction sounds crystal-clear at first, but in reality it is mud-blurry. Here is some of what Flash wrote me:

For a sentence to attempt to say what it shows is to commit an error of logical types. It seems to be putting a round peg into a square hole, whereas it is instead putting a round peg into something which is not a hole at all, square or otherwise. This is a category mismatch, not a paradox. It is like throwing the recipe in with the flour and butter and eggs. The source of the equivocation is an illegitimate use of the term ‘this’. ‘This’ can point to virtually anything, but ‘this’ cannot point to itself. If you stick out your index finger, you can point to virtually anything; and by curling it you can even point to the pointing finger; but you cannot point to pointing. Pointing is of a higher logical type than the thing which is doing the pointing. Similarly, the referent of ‘this sentence’ can be virtually anything but that sentence. Sentences of the form exemplified by ‘This sentence no verb.’ and ‘This sentence has a verb.’ are not well-formed: they commit fallacies of logical type equivocation. Thus their self-referential character is not genuine and they present no problem as paradoxes.

There will always be people around who will object in this manner, and in the Brabnerian manner. Such people think it is possible to draw a sharp line between attributes of a printed sentence that can be considered part of its form (e.g., the typeface it is printed in, the number of words it contains, and so on), and attributes that can be considered part of its content (i.e., the things and events and relationships that it refers to).

Now, I am used to thinking about language in terms of how to get a machine to deal with it, since I look at the human brain as a very complex machine that can handle language (and many other things as well). Machines, in trying to make sense of sentences, have access to nothing more than the form of such sentences. The content, if it is to be accessible to a machine, has to be derived, extracted, constructed, or created somehow from the sentence’s physical structure, together with other knowledge and programs already available to the machine.

When very simple processing is used to operate on a sentence, it is convenient to label the information thus obtained “syntactic”. For instance, it is clearly a syntactic fact about “This sentence no verb.” that it contains six vowels. The vowel-consonant distinction is obviously a typographical one, and typographical facts are considered superficial and syntactic. But there is a problem here. With different depths of processing, aspects of different degrees of “semanticity” may be detected.

Consider, for example, the sentence “Mary was sick yesterday.” Let’s call it Sentence M. Listed below are the results of seven different degrees of processing of Sentence M by a hypothetical machine, using increasingly sophisticated programs and increasingly large knowledge bases. You should think of them as being English translations, for your convenience, of computational structures inside the machine that it can act on and use fluently.

1. Sentence M contains twenty characters.
2. Sentence M contains four English words.
3. Sentence M contains one proper noun, one adverb, in that order.
4. Sentence M contains one human's name, one linking verb, one adjective describing a potential health state of a living being, and one temporal adverb, in that order.
5. The subject of Sentence M is a pointer to an individual named 'Mary', the predicate is an ascription of ill health to the individual so indicated, on the day preceding the statement's utterance.
6. Sentence M asserts that the health of an individual named 'Mary' was not good the day before today.
7. Sentence M says that Mary was sick yesterday. one verb, one adjective, and

Just where is the boundary line that says, "You can't do that much processing!"? A machine that could go as far as version 7 would have actually understood-at least in some rudimentary sense-the content of Sentence M. Work by artificial-intelligence researchers in the field of natural language understanding has produced some very impressive results along these lines, considerably more sophisticated than what is shown here. Stories can be "read" and "understood", at least to the extent that certain kinds of questions can be answered by the machine when it is probed for its understanding. Such questions can involve information not explicitly in the story itself, and yet the machine can fill in the missing information and answer the question.

I am making this seeming digression on the processing of language by computers because intelligent people like Mr Flash qFiasco seem to have failed to recognize that the boundary line between form and content is as blurry as that between blue and green, or between human and ape. This comparison is not made lightly. Humans are supposedly able to get at the "content" of utterances, being genuine language-users, while apes are not. But ape-language research clearly shows that there is some kind of in-between world, where a certain degree of content can be retrieved by a being with reduced mental capacity. If mental capacity is equated with potential processing depth, then it is obvious why it makes no sense to draw an arbitrary boundary line between the form and the content of a sentence. Form blurs into content as processing depth increases. Or, as I have always liked to say, "Content is just fancy form." By this I mean, of course, that "content" is just a shorthand way of saying "form as perceived by a very fancy apparatus capable of making complex and subtle distinctions and abstractions and connections to prior concepts".

Flash qFiasco's down-home, commonsense distinction between form and content breaks down swiftly, when analyzed. His charming image of someone making a "category error" by throwing a recipe in with the flour and butter and eggs reveals that he has never had Recipe Cake. This is a delicious cake whose batter is made out of cake recipes (if you use pie recipes, it won't taste nearly as good). The best results are had if the recipes are printed in French, in Baskerville Roman. A preponderance of accents aigus lends a deliciously piquant aroma to the cake. My recommendation to Brabner and qFiasco is: "Let them eat recipes."

reprinted at <http://www.IniciativaAtea.org/foro/viewtopic.php?f=26&t=3992>

[End of Mr. Hofstadter's passage; the bit quoted as mine was excerpted from correspondence during the time when he edited the Mathematical Games column at *Scientific American*, a column he edited with great acumen and which I regularly read with interest. --FqF]

Reply to Mr. Hofstadter:

The transition between day and night is not a boundary line, but a period of twilight. This does not mean that we cannot distinguish day from night. Similarly, the fact that there is no border between blue and green does not mean that blue can or should be assimilated to green. Similarly, the fact that a few borderline cases can be conjured up which blur the distinction between some forms and some contents does not mean that we cannot or should not continue to distinguish form from content generally. The distinction is much too obvious, useful, and powerful to be collapsed into one of its terms generally, a few blurry cases notwithstanding.

I should clarify two points at the outset. First, the form of a sentence is not restricted solely to its printed symbols or its spoken sounds. Any sentence may exhibit many forms depending on the aspect under which it is to be considered. Second, sentences and statements should be distinguished, no less so than letters and fonts and alphabets should be. Letters, sentences, and statements may all take many different forms, and *types* of forms, concurrently; so it is misleading to talk about *the* form of a sentence, as if it could have only one. For example, letters may belong to one or another alphabet (Latin, Cyrillic, Hebrew, Arabic, etc.); that is one type of form. Within an alphabet, the font would be an example of another type of form (Times New Roman, Arial, Garamond, etc.). Within a font, there may be other types of formal difference or formal similarity (serif, sans serif, italic, bold, etc.). There may be any number of different forms within any given type; similar layers of multiplicity apply to sentences and statements as well.

The difference between a sentence and a statement is this: a sentence consists of letters and is therefore specific to an alphabet; a *statement* is the assertion of a proposition with a sense (independent of any specific alphabet). A sentence and/or statement may have many forms, not necessarily only when in print, concurrently. A sentence may, for example, be grammatically well- or ill-formed (syntax). Other examples of forms include: assertion, question, command, performative (“you’re fired!”, “I sentence you to 30 years hard labor,” “you are under arrest”), exclamation, imprecation (“expletive deleted”), exhortation, accusation, operative (“open sesame”), etc. etc.; poetry, prose, prosody, Gregorian chant; a rule of grammar, an example of correct usage (to teach someone), an instance of correct usage--to indicate but a few.

Hofstadter’s examples of machine processing of various sentences do not show that form is fancy content, as he claims. What his examples show is that humans are very clever computer programmers. Form and content are still distinguishable, even here; what his examples show is that one can choose to ignore certain distinctions for certain purposes within certain rigidly specified parameters. The difference is still there though, even when you ignore it.

I fully accept that there are different degrees of semanticity; that still does not prove that everything linguistic can or should be reduced to semantics. Grammar is not vocabulary, or vice versa. A rule of the game should not be assimilated to a move in the game, or vice versa. I do not believe that Hofstadter believes that blue can and should be assimilated to green, *because* there is no borderline between blue and green. So his analogy between blue and green, and form and content, does not hold.

There are cases in which form and content really cannot *not* be distinguished; disregard of the distinction, or collapsing the distinction to “fancy content”, is untenable. For example, in dactyl hexameter [dum-diddy / dum-diddy / dum-diddy / dum-diddy / dum-diddy / dum-dum] or Iambic pentameter [ba BUM / ba BUM / ba BUM / ba BUM / ba BUM]. Ignore the form and you have missed the whole ball of wax. The fact that form can be represented by nonsensical placeholders (dum-diddy) is the proof that form is not fancy content.

Or, to take a more pedestrian example:

*“This is rhyming verse,
However terse.”*

Whether that is poetry or not pertains to the form, not the content (sense). Collapse the one into other and you are well along the high road to incoherence.

Similarly, distinguishing prose from poetry is a matter of form, not content (sense). This is not to say, however, that there is a clear borderline between every instance of prose and every instance of poetry (“*If called by a panther, don’t anther.*” --Ogden Nash). There need not be a clear borderline; as Wittgenstein said, family resemblances and overlapping cases may come into play, without vitiating the general validity of the distinction.

Form and content are not to be applied rigidly, the same in every case. What is form in one context or level of discourse may be content in another. What counts as form for one criterion of distinction or of comparison may count as content for another. This flexibility is what makes the distinction so very useful and powerful.

Imagine someone saying “look at the difference between these two chairs.” And then the person addressed looks *between* the chairs. That is the sort of nonsense which results from willful disregard of categories.

Hofstadter and I share a joy in paradoxes. However, where I take joy in genuine paradoxes, he has a tendency to sometimes impute paradoxicalness to things which are merely incongruous or muddled. The logic of self-referential statements is a case in point.

“This sentence no verb,” for example. In this case, the content and the form seem to collide, the one apparently negating the other. Correctly understood, however, there is no paradox here and indeed the collision is merely apparent. What we have here is not a complete sentence, for it is syntactically malformed; therefore, in the above string (to use his favored idiom of informatics) “this sentence” is not self-referential, since, grammatically speaking, there is no sentence there to refer to. Whether a sentence is syntactically correct or not depends on its conformity or lack thereof to grammatical rules; that is, it depends on *structure*, not content. This is the point about the claim that a sentence cannot say (content) what it shows (structure). The fact that one can contrive sentences which appear to conflate structure (or form) with content--“this sentence contains five words”-- is no counter-example, for such sentences still have higher-order structures which are *not* said within the sentence (for example, whether the sentence is syntactically correct). “This sentence is syntactically correct” also has higher-order forms and structures not said (for example, that it is in English and not German). And so on and so on.

What we have in the case of “this sentence no verb” is pseudo-self-reference within a pseudo-sentence. There is just barely enough content (sense) for a human to find it mildly amusing or puzzling, by surreptitiously supplying the missing verb. Change the word order, however, and all seeming paradoxicalness evaporates: “this verb no sentence,” “this no sentence verb,” “no sentence verb this,” “no this verb sentence,” etc. In the first of those gibberish strings--“this verb no sentence”--, it is obvious that “this verb” designates no verb at all (since “verb” is a noun); analogously to “This sentence” in “This sentence no verb.” Hereagain, the distinction between something which almost makes sense, “this sentence no verb,”--that is, something which can be made into a sentence by a plausible addition or substitution--and complete gibberish, “this verb no sentence,” need not be a borderline in order to be useful and intelligible (“*All mimsy were the borogroves, and the mome raths outgrabe.*” --Lewis Carroll).

“This statement is false,” is no paradox for a similar reason. There is no statement there, so there is nothing to *be* either true or false. There must be some more propositional content than merely “this statement” in order for it to be capable of bearing a truth value. “This statement is false” is analogous to “what is the answer to this question?” There isn’t a question there, despite the “?” at the end of it. This becomes clearer if we formulate it thus: “Is this number odd or even?” “This number” does not designate any number at all, so oddness or evenness is moot. The same reasoning applies to the following pair: “A: B is true.” “B: A is false.”--as well as to all analogous puzzles of the form “The statement on the back side of this card is true” “The statement on the front side of this card is false.” The problem is not that “this card” is self-referential (which it

isn't--cards do not refer) but rather that "the statement on the back/front side" does not have enough content to be true or false. So how much content is enough? How long is a piece of string?

Alternatively, one may consider this in light of Frank Ramsey's insight that "p is true" is logically equivalent to "p"--the "is true" adds nothing. Thus, "This statement is true" reduces to "This statement." And now we may legitimately demand, *what* statement? There just isn't any statement there, and tacking "is true" onto the stub doesn't make it into one. Analogously, "p is false" reduces to "not p", and "This statement is false" to "Not this statement." Since "this statement" states nothing, "not this statement" negates nothing. That is no paradox.

There are cases in which self-referential sentences are perfectly well-formed and understandable. For example, in typography:

"This is bold italic."

"This is 24 point Arial."

Such examples present no problem logically and represent no paradox of self-referentiality. They are simply *examples of fonts*. Neither is there anything paradoxical about the following example:

"This is 128 point Arial bold italic."

It just isn't (128 point or Arial or bold or italic); so the statement is simply false, not paradoxical.

I do not know how to read Mr. Hofstadter's comment that my "... charming image of someone making a "category error" by throwing a recipe in with the flour and butter and eggs reveals that he has never had Recipe Cake." Well, of course, I have never had Recipe Cake and neither has anyone else, I'll wager. It would taste pretty flat.

The concept of category error is not mine. Nor is it "*down-home*," if by that Mr. Hofstadter means parochial. It is a commonplace among logicians to preserve category distinctions; failure to do so is known as committing a "howler", like confusing the map for the territory or the menu for the meal. It is surprising if Mr. Hofstadter thinks he can dismiss such weighty matters with the virtual wave of a computer hand. A computer program is an artificial language, and artificial languages, like natural ones, obey certain fundamental rules and distinctions--they too exhibit grammar (schematic forms) and specific content which instantiates the grammatical forms; there is nothing to be gained by confusing artificial form and content and very good reason not to.

Someone once said that a story has a beginning, a middle, and an end. So, is the following a story? "This is the beginning. This is the middle. The end." No. That is not a story, despite having a beginning, a middle, and an end. It is a schema for a possible story, and the sentences comprising it are placeholders where the content would go, if it were a story.

Logicians are familiar with similar schemata: "p implies q, p, therefore q." That is not an argument, for it is devoid of content--it is an argument *form*; p and q are placeholders, representing some possible content. A reduction of form to content is not possible in such a case.

David Moser wrote:

This Is the Title of This Story, Which Is Also Found Several Times in the Story Itself

This is the first sentence of this story. This is the second sentence. This is the title of this story, which is also found several times in the story itself. This sentence is questioning the intrinsic value of the first two sentences. This sentence is to inform you, in case you haven't already realized it, that this is a self-referential story, that is, a story containing sentences that refer to their own structure and function. This is a sentence that provides an ending to the first paragraph.

This is the first sentence of a new paragraph in a self-referential story. This sentence is introducing you to the protagonist of the story, a young boy named Billy. This sentence is telling you that Billy is blond and blue-eyed and American and twelve years old and strangling his

mother. This sentence comments on the awkward nature of the self-referential narrative form while recognizing the strange and playful detachment it affords the writer. As if illustrating the point made by the last sentence, this sentence reminds us, with no trace of facetiousness, that children are a precious gift from God and that the world is a better place when graced by the unique joys and delights they bring to it.

This sentence describes Billy's mother's bulging eyes and protruding tongue and makes reference to the unpleasant choking and gagging noises she's making. [and so on]

I hate to disappoint Mr Moser, but there is nothing self-referential about his so-called story, even though each of the sentences is self-referential. Moser commits the same mistake as Hofstadter, confusing content for form. Self-referential statements do not a self-referential story make. What they make is a list of self-referential statements. A self-referential story would be one in which the composing and the telling or writing of the story *were* the story. Such a story might, for example, begin "I, ... Claudius.... am now about to write this strange history of my life..." during the course of which it is revealed that the writing of the story materially affects the course of his life and thereby also the content of the story.

Another self-referential story might begin and end: " I have this terrible writer's block. *fin*"

fin