

# Notes on Holbo's drafty beachcombing

Tikitu de Jager

October 2007

## Abstract

This is an attempt to apply (my understanding of) Ruth Millikan's theory of functions to John Holbo's extension of Knapp and Michael's paradox of wave poetry, aiming at a theory of meaning that's applicable to literary criticism. I've not much background in literary criticism (you will have noticed that this is an "abstract" not an "introduction"), and the notions that strike me as naïve are doubtless more firmly grounded than it seems to me. Please try to read my ignorance charitably rather than as arrogance, although of course I don't expect any such restraint regarding my argument.

Looking for meaning is tricky. I'm going to start somewhere a bit simpler, with the question of whether the shapes Smith writes on the beach are really letters forming words forming a poem. Once I've argued that, *contra* K&M, in a relevant sense they are, then we can get to the question of what they mean.

In Millikan's story, our concepts are useful (and hence reproduced, by learning and explicit communicative transfer and so on) in so far as they track objectively existing "clumps, humps and peaks" in the space of logical possibility. These clumps group objects that are similar for causal reasons, so that our concepts support inductions (tomorrow's dog is similar to today's not by coincidence; we can learn something about dogs by observing a single dog *because* the resemblance between dogs is not coincidental).

Under this story, Smith is not a person; or, to be very careful: he doesn't belong to the same ontological kind as Wordsworth. Millikan's objection to the swampman thought experiment is that my first formulation contains another source of complication: we have a word "person" and a corresponding concept, and it's entirely unclear what we should do with these if Smith exists. If there's something that behaves entirely like Wordsworth but *by coincidence* then having a concept of "Wordsworth" might not even be coherent. (This is the essence of the Davidson quote that Brad posted. Millikan distinguishes an ontological level underlying our means of identifying objects [these might not be quite what Davidson means by 'criteria' since they're fallible, but the notion is similar] but this level only has primacy inasmuch as it behaves regularly. To introduce Smith is explicitly to undercut the regularities that support the notions the example is trying to probe.)

This can't be quite the entire story though. We can see this by dialing back the unlikelihood of the coincidental events, and similarly restricting the grandeur of the interpretations being applied. Similarly to Holbo's typographical variant of the K&M story, this will suggest that the notion of "functional

kind” might make sense at an ontological level, without requiring the sort of causal connections that Millikan requires. There may be categories of things that support reidentification and induction primarily via the behaviour of interpreters, and only contingently via causal connections. (If there are such animals, the interesting question will be whether poems might be so described, and whether the idea of meaning is compatible with such a picture.)

## 1 Found typography

Let’s take, instead of poems of Wordsworth, individual letters as our found objects. It’s easy to imagine finding a wave-tossed clump of seaweed that, when squinted at the right way, looks like an A. Through diligent beach-combing over many years, John Holbo might collect a complete alphabet of photographs of these letters, and even use (copies of) them to typeset a poem (probably, respecting tradition, “A slumber did my spirit seal”).

According to a strict Millikanesque reading, the photographs he takes *are* letters. They have a causal connection to the more ordinary letters appearing on your screen or paper, albeit a more complicated one than normal. A normal causal connection between two occurrences of the letter ‘A’ is that one was copied from the other, or that both descend historically from some common ancestor. The connection here is more twisty: the photographs Holbo takes are *chosen* because of a similarity he perceives between a printed letter and the seaweed. They are in a sense copies of printed letters, but reproduced via a complicated process involving Holbo’s capacity for recognising similarities and his beach-combing and selection of good photographic possibilities. (To see that there is still something like copying going on, suppose that our alphabet was different; the photographs Holbo chose to make would differ in systematic ways, matching the systematic differences between alphabets.)

So far so good, Holbo’s photographs are letters (in the sense that they belong in the same ontological kind as normal letters: they are the way they are *for the same reasons* normal letters are the way they are). But what about the original seaweed configurations? According to Millikan, these *don’t* belong in the same ontological category.

Now, though, let’s make this problematic. Because as well as a photographer, John Holbo is a secret installation typographer. Over the years he has photographed many hundreds of found letters, yes, but he has also *dug them up*, with painstaking care to avoid damaging their configurations, and placed them in planter boxes in his back garden. He’s using them to write a poem.

Now are they letters?

Certainly the words they form are words; for they stand in the same complicated causal relation to ordinary words as the photographs do to ordinary letters. (Assuming that the letters are found individually, none of the words in Holbo’s back garden would exist without a normal word for it to be copied from, although again the copying is by a rather complicated mechanism.) And there is the overwhelming impulse to attribute Cummins-function to the letters that make up the words, because clearly (?) they wouldn’t be words if their components weren’t letters.

But now suppose that Holbo’s poem is completed, but he goes on beach-combing. Every now and then he finds a better A, or a more shapely S, and

replaces one of the pots in his matrix with the new find. And he returns the rejected letter to the beach, depositing it entire on the sand and leaving it there.

Is it still a letter?

## 2 Intentionality

There is a sense in which the seaweed really changes character just by being placed in a beach context or an array of planter pots. In Holbo's back yard we can say "that's not a very good E", whereas on the beach we have to say "that doesn't look much *like* an E". We can make normative judgements about the properties of an object when those properties are responsible (in part) for the existence of the object. Millikan's emphasis on copying is partly a consequence of her concern with biology, but partly also a response to counter-examples like the two rocks in the stream;<sup>1</sup> the point of the found typography example is to show that the right kind of selection, without copying, can still lead us to make the same sorts of normative judgements. The question is of course whether we can ground these in some way, as concretely as Millikan does through the notion of copying.

It seems to me that we can do so by reversing the temporal direction the definition looks in. That is, a first attempt is to say that the seaweed in Holbo's garden is a letter because it will be copied in a way that creates letters, while the seaweed on the seashore is not a letter *unless Holbo is destined to dig it out*. (Is Smith's poem — or worse still the entirely coincidental shapes eroded directly by the waves — are they still a poem if it so happens that nobody walks past before the waves erase them again?)

Yes, this means we can't know in advance and in a theoretical sense whether the wave poem is a poem. Holbo has said on *The Valve*, "We, the experiment's audience, know something that no one in the situation actually would know"; under my story, it's this perspective that makes it impossible to call the wave poem a poem. We're not actually seeing it, we're asking "what if it existed?" and we're not saying yet what would happen in its future. But if anyone in fact comes upon it on the beach, if they speak English they *will* copy the letters and words it contains in the very act of reading. (I'm thinking firstly of a "neural copy", but if this raises philosophical issues for you then imagine the finder writing the poem into their pocket notebook, or reciting it to a friend when they get home, or whatever. The existence of these copies is more contingent than the neural one — suppose the finder gets eaten by a sea-lion on his way home — and I like the notion that positing an interpreter necessarily makes the

---

<sup>1</sup>Quoting Holbo on Wright,

The function of  $X$  is  $Z$  means

1.  $X$  is there because it does  $Z$ ,
2.  $Z$  is a consequence (or result) of  $X$  being there.

The rocks counterexample takes the form of two rocks in a fast-flowing stream, one sitting on top of the other. The lower rock is there because it supports the higher (otherwise the stream would sweep it away) and it does so because it is there, however we don't want to say that supporting the other rock is its function. Millikan's move is, roughly, to require that the first clause make reference to copying, so that  $X$  was copied from some  $Y$  because  $Y$  did  $Z$ . With some further tightening of definitions we can get that  $X$  *doesn't even perform*  $Z$ , so that the notion of a defective object that does not fulfill its function becomes grounded in causality.

poem into a poem, but there are much more interesting points to argue about than the ontological status of neurophysiological states.)

### 3 Oops...indeterminism

Now it seems we have a problem though. The poem-hood of our poem seems to be dependant on whether the reader is literate, for a start. If we try to extend our account to the meanings of the words, we'll find that they mean different things depending on the language the interpreter speaks. In other words, this story succumbs to radical indeterminism.

Well, yes and no. It seems to me that a notion of *the* determinate meaning of a linguistic expression (token) is simply incoherent, talk of authorial intention notwithstanding. I don't even mean that such a notion is possible in the abstract, that meanings-as-intentions have ontological existence but we can never access them or know that we have discerned them. Determinate meanings may exist; authorial intentions certainly do. What I'm rejecting is the requirement that linguistic expressions are necessarily paired with them.

(Here, if anywhere, I expect to be given a good spanking for making stark claims that have been chewed over and discarded years ago in the critical literature.)

The point is that even Millikan's own copying theory of ontological kinds, or an authorial-intention theory of meaning, can be tweaked to produce indeterminacy. It's possible for an utterance (token) to be copied from *multiple* previous tokens; it's possible for the same utterance (token) to be the result of multiple intentions.

The Valve comments on collaborative composition inspired me in this direction, but I think you still need to do some extra work to make them problematic. Specifically, it's not enough to say that multiple authors are partially responsible for the work as it stands and that it's not possible in practise to attribute intention; this leaves the possibility of a true ontological-level meaning which we are simply practically incapable of uncovering. Imagine though that we (you and I, for a moment) are collaborating on a poem, and both independently suggest the opening line "A slumber did my spirit steal". It so happens that we had different reasons for the choice, but we might agree to leave the line as it stands, intentionally ambiguous (ambiguity intentional). Suppose though that our explanatory letters get lost in the post, so although the poem is published with that first line neither of us realises that the other intended a different reading.

The point of this excess is that there is not, even in principle, any way to divide the line into components and attribute unique intentions per component. Nor is there any intention (if the post didn't get through) to write ambiguously.

With similar effort (and similar degrees of likelihood) we can play the same trick on a single poet, if he is sufficiently forgetful. We should not, I think, be put off by how unlikely and forced these examples are. The past is fixed while the future is fluid, so it's not surprising that "forward-looking" function attributions are less determinate than "backward-looking" ones; furthermore, the paradigm cases of linguistic expressions are reproduced *because* they give good, stable clues to the intentions that lead to their production. Certainly the unlikelihood is not in the same class as the production of a Smith or a

wave poem, and doesn't lead inexorably to the sort of conceptual breakdown or rearrangement that these do.

So what we get, in the end, is more or less a common-sense picture of meaning. The conventional meaning of a word, as we readers understand it, is a good clue to the meaning of the same word in a poem; our knowledge of the author's intention (if we can get this from any other source) is another good clue; a third clue is our personal reaction, as evidence of what factors might lead to our reusing the word ourselves (my forward-looking addition to the story). If we're lucky these will be more or less aligned and we can talk about "the" meaning of the word; if we have good evidence that they *don't* align, though, none of the three has any special ontological status. Our concept of "meaning" is confused in Millikan's terms (doesn't track a kind with causal connections), and with some effort we can give evidence of this by showing cases where induction on meanings doesn't succeed.<sup>2</sup>

#### 4 (The same goes for texts)

At the level of texts, if one wants to say that a text has a single determinate meaning, the counterexamples get much easier to produce. In a collaborative project it can easily happen that different individuals intend the work as a whole to have a different interpretation, and so on. I've started instead with words because you can probably duck around the problem, or at least expend a lot of effort trying, when you can chop the text into pieces and try to assign responsibility piecewise.

We also need the same sort of forward-looking notion as for the found typography to cope with things like found poetry, computer-generated art, and so on. It's a partial answer to say that the programmer is the author of a randomly generated poem, but the individual lines have meaning also by virtue of the conventional meanings of the words they are composed of (these words are words in virtue of having been copied from other words, written into a code dictionary and so on) and also in virtue of their future interpretation as meaningful sentences. A population of computers swapping poems among themselves would be writing words, and possibly sentences, but not poems ... until someone read them, and published an anthology of the best ones.

---

<sup>2</sup>In a sense all that K&M do is respond to these examples by throwing them out of the extension of "meaning"; this is going to lead to unconfused concepts in Millikan's terms, but terribly confused communication with people who associate the same words with different concepts.