

On Being Inexplicit.

JAKE CHANDLER

School of Cognitive and Computing Sciences, University of Sussex, Falmer, Brighton BN1 9RH, United Kingdom. (e-mail: jakec@cogs.ac.uk)

CSRP 516 - March 2000.

Abstract: Dennett's singular position on the status of beliefs and desires can be characterized by a negative claim: beliefs and desires are not necessarily internal states involved in the aetiology of behaviour. Motivating this claim is the recognition of a class of belief / desire assignments in which there is no explicit representation tokened in the system: we are said to be dealing with 'inexplicit', or 'tacit' representation. But what exactly is 'tacit' representation? The problem is to find a naturalistic alternative to the account of beliefs and desires as internal content-bearing states, which will embrace this class of inexplicit representation, both supporting univocal assignments and granting these assignments explanatory bite. While everyone is familiar with Dennett's 'Intentional Stance' story, an alternative position is found to be compatible with, and indeed even suggested by his writings. An appeal to biological teleology is made, normal conditions for proper functioning of behaviour being said to be tacitly believed (in case there is no internal state purporting to coordinate behaviour with the presence or absence of those conditions), and normal telic outcomes of behaviour said to be tacitly desired (again in case there is no internal state purporting to elicit behaviour in the presence of appropriate conditions). The concepts of belief and desire are cast in terms of relational properties holding between an organism and states of affairs by virtue of the possession of a trait with a certain function.

Key words: Belief / Desire Psychology, Dennett, Explicit, Fodor, Implicit, Representation, Teleology, Tacit.

1. Fodor on Beliefs and Desires.

Jerry Fodor claims that cognitive science will vindicate folk psychology through the following pair of implications:

"For each tokening of a propositional attitude, there is a tokening of a corresponding relation between an organism and a mental representation".

"For each tokening of that relation, there is a corresponding tokening of a propositional attitude".

(Fodor 1987:20)

What does this mean? First that folk psychologists, you, me, or my next-door neighbour, have reasonably strong intuitions concerning the ontological nature of beliefs and desires. The folk tools come with a folk manual and its folk glosses. A minimum commitment is thought to be made to the existence of discrete internal states, which have semantic content, and which make a causal contribution in the aetiology of behaviour. This much Fodor finds vulgar to squabble about, says he.

Not only is this what we think beliefs and desires are, but this picture is what Fodor bets a mature cognitive science is going to back up, with a few technical refinements thrown in. Fodor thinks for instance that the classical computer metaphor of the mind is suitable, with mental representations as syntactically well-formed concatenations of symbols with a combinatorial semantics, and that the 'attitude' aspect of propositional attitudes will be rendered by a causal-functional role of the sentence token, the fact, as Fodor puts it, that mental representations occur in "belief boxes" or "desires boxes". These details are however specific to Fodor's particular views concerning the proper treatment of a scientific vindication of folk psychology. They are only secondary to his more general claim that what the terms belief or desire routinely pick out are internal representations of some sort, and that this is what we think they pick out.

Now this very point is of course what Daniel Dennett has problems with. As he himself puts it "realism with regards to beliefs as 'discrete internal states'... has been my chief stalking horse" (Dennett pers. comm. in Bechtel 1985). Indeed, leaving aside what Dennett thinks we think beliefs and desires are, much of Dennett's work is peppered with examples purporting to show that we quite frequently and quite naturally fail to ground our belief / desire attributions in existent thusly semantically interpretable internal states. Talk is of "potentially explicit", "implicit", "emergent", or again "tacit" representation is popular here, though, rather frustratingly however, it must be said that the discussion does tend to stick to the anecdotal.

Fodor of course isn't impressed by the examples given. He has hedged his claims to the effect that belief / desire attributions pick out relations to mental representations, ... except when they don't. Basically, picking out internal, contentful, etc, states is what we should be doing. It is when, and only when, this happens that our attributions have explanatory bite.

This point deserves some attention. The idea is that, in many scientific co-optations of everyday terms, it turns out that the bearer of the properties crucial to our adoption and use of the term is in fact a member of a much narrower class than the one we were picking out prior to scientific investigation. We had overextended the domain of application of the t

erm. Fodor's example here is water: we used to call a whole range of things "water", but now we know better, or at least we know when to temper our qualifications with e.g. 'chemically impure'. The real bearer of those properties which, for example, enabled the whole thirst-quenching business, is H₂O. The "proper" applications of the term water thus involves chemically pure samples of H₂O. These are the "core cases", the rest are somehow "derivative".

Analogously, according to Fodor, what is really enabling our belief / desire based explanations to do the job they do is the presence of these internal, contentful, causally potent states in the system whose behaviour is being explained. Whatever does not refer to such things is attribution manquee, and can have no explanatory relevance.

Fodor does maintain however that for the core cases, explicit representation is crucial. That is, if one were to somehow show that, for a given belief / desire attribution, some thusly contentful internal state could be seen to be picked out, and that content was not explicitly represented, he would be in trouble. "No Intentional Causation without Explicit Representation" claims Fodor (Fodor 1987: 25). I'm not quite clear here as to why he exposes himself to such a line of criticism, presumably because of his additional commitments to Language of Thought, but whatever the motives, note then that here is room for a first swing.

The second possible move would be a challenging of the choice of core cases., i.e. the idea that there can be No Intentional Explanation without Intentional Causation, as Fodor might put it. This is the kind of move Dennett would need to make. In order to satisfy Fodor, his response here would have to be one that shows that his cases of non internally tokened belief / desire attributions can, without then invoking internal causes, have scientific explanatory credentials. This Fodor very much doubts, as the only other realist option he sees is his one, save perhaps some kind of behaviourism, which as everybody knows, for all else that's wrong with it, can't be put to any explanatory work.

Such scepticism is of course heir to a now long and well established tradition stemming from Davidson's "Actions, Reasons and Causes" (Davidson 1963), the central argument of which took the form of a challenge: what on earth could provide explanatory bite to belief-desire attributions save a causal story (within the realm of the naturalistically kosher of course)?

"failing a satisfactory alternative, the best argument for [the causal theory] is that it alone promises to give an account of the 'mysterious connection' between reasons and actions." (Davidson 1963: 11)

Orthodoxy no wants it that the causal story is the only serious game in town. But is it? Whilst perhaps not articulating Dennett's most 'official' take on the issue, I take a close look, in the next section at the cases of inexplicit beliefs / desires at the root of his hostility towards the 'inner cause' view, take a fine-tooth comb to pick out some relevant statements in Dennett's writings and find that a naturalistic alternative strongly suggests itself.

2. Dennett on Beliefs and Desires.

Here's Dennett on inexplicit representation then.

First of all, a brief comment on what counts as explicit representation. Dennett himself gives a succinct definition in 'Styles of Mental Representation' (Dennett 1982): "information is represented explicitly in the system if and only if there actually exists in the functionally relevant place in the system a physically structured object, a formula or string or tokening of some members of a system (or 'language') or elements for which there is a semantics or interpretation, and a provision (a mechanism of some sort) for reading or parsing the formula" (emphases omitted). Note that a generally widespread concern is to rule out from the definition things like the weighting of connections between nodes in a connectionist network. There is a clear feeling that whatever information there is there cannot be subject to judgments concerning factivity or be involved in further reasoning. It is not information for the system one would say.

David Kirsch (1991) would no doubt worry here about grounding explicitness in what he calls a "structural" definition, one that doesn't appeal to considerations of processing time, of computational complexity required to make the information readily usable. He feels very strongly about denying the status of explicit information to information that takes time and resources to extract. For him the crucial upper limit is constant time access, longer than that and the information is inexplicit. In particular, the parsing and semantic interpretation of first order predicate calculus, thought to be a good candidate for being an analogue, complexity-wise, of LOT, is a non-constant-time process. I don't think this matters here however: once the information is extracted and put to use, it is explicitly represented. Let us call this inexplicit information 'potentially explicit' and see whether Fodor can deal with it or not.¹

It's time then to get out of the way then the objection that the beliefs and desires typically predicated to us may be simply potentially explicit, i.e. internal contentful causally potent states, potentially derivable, either in principle or simply in practice, within the resources of the system from other such states (Dennett 1979a:104-106; 1981; 1982). The problem with this kind of rebuttal here is of course Fodor's readiness to make the kind of dispositionalist concessions that Dennett perhaps wasn't expecting of him, whilst of course maintaining that, though they can be attributed, dispositional beliefs cannot play an explanatory role vis a vis behaviour (Fodor 1987:22)

Another type of inexplicit representation invoked here is 'implicit' representation, defined as what is implied logically by something that is stored explicitly (Dennett 1982). I'm not quite sure however, that Dennett has shown just quite how this implicit information can have explanatory relevance short of being potentially explicit. His 'murder in Trafalgar Square' example (Dennett 1981), which presumably purports to show how a proposition logically implied by reasonably different explicit informational states in four different people can ha

ve predictive or explanatory value (e.g. predicting similar behaviours in the face of a question concerning the factuality of that proposition), seems to me a simple case of dispositional belief. The four men who implicitly believe that P behave the same way when faced with the question 'is it the case that P' simply because they can actually derive the logically implied proposition, i.e. render it explicit, when adequately prompted. The justification for the relevance of implicit representation in "Styles of Mental Representation" isn't crystal clear either. I suggest we suspend judgement until a clear example is provided and defended.

The contention that we ground, e.g., belief attributions in implicit and potentially explicit information isn't however the main thrust of Dennett's arguments. The real key notion is the idea of tacit representation. The main idea is that tacit representation is representation that is never tokened in the system, nor is implied by explicit tokenings but nevertheless plays an explanatory role. As we shall see however, Dennett's account is found guilty of being rather woolly on two fronts (1) he gestures in the directions of many prima facie different types of candidates for tacit representation-hood without attempting to provide a systematic account, in his own words, "the critical term, 'tacit', still has been given only an impressionistic, ostensive definition" (Dennett 1982) (2) more seriously, he doesn't make clear what the attitude relations between these tacit representations and the system that represents them are, he tends not to explicitly qualify them in terms of beliefs or desires.

The main examples discussed here involve tacit representations of rules, like the following of the rules of arithmetic by a pocket calculator, or, though not discussed by Dennett, but faithful to the spirit here, the following of Ohm's Law by a connectionist network (as described in Smolensky 1988). Nowhere is there any kind of symbolic representation that the system consults. The system is said to 'honour' the rule without explicitly representing it. What should one say in connection to belief and desire here? Dennett doesn't say anything himself. Perhaps we should say that the system knows the rule, or again that it believes it to be true. Smolensky's network would then believe that "it is the case that $V = C \cdot R$ ".

Now of course, the rules that are said to be honoured don't necessarily correspond to the actual "laws of thought" mediating between the explicit representations of, say, input and output. Indeed, to put things the other way round, the latter may only roughly approximate the former (as would be expected for e.g. a network trained on a limited number of exemplars, i.e. legal combinations of values). Why then do we not invoke tacit knowledge of those other, actual, relations mapping inputs to outputs? The answer seems to hinge on the fact that following the rules of arithmetic or Ohm's rule is a competence characterisation. The function of the network is to produce outputs for given inputs consistently with the given laws. Were those laws to be different, the mappings between the explicit beliefs would have to change accordingly. The truth of the laws is a crucial environmental condition for the proper functioning of system2. We are now equipped to suggest a first account of tacit belief:

(A) An organism O which:

(i) has a trait T which has a function to bring about E by means of doing R, and this relying on circumstances A_1, \dots, A_i to perform satisfactorily, and

(ii) does not have a trait (a) the function of which is to adapt T by eliciting T in the sole presence of circumstances A_1, \dots, A_i , and (b) which is in a state such that it itself relies on the obtention of circumstances A_1, \dots, A_i for proper functioning.

is said to tacitly believe that A, B and C obtain. It is said to explicitly believe circumstances A_1, \dots, A_i obtain in case (ii) fails.

This tentative characterisation does seem to make sense in the light of some of Dennett's other comments. Indeed, elsewhere, he has alluded to something that seemed prima facie quite different to calculator-tacit-arithmetic-rule-following, more in the spirit of what Millikan calls 'tacit presuppositions' (Millikan 1993:104-105). In "True Believers", he talks of "representation of the environment in - or implicit in - the organisation of the system", as "a way of alluding to this tight relationship that can exist between the organisation of a system and its environment". Given behaviours or features of the system demand "very specific environment[s] in which to operate properly" (Dennett 1979).

The role of the assumptions made by an organism, when carrying out a piece of behaviour, in fixing the content of a triggering indicator, has eminent support and equally eminent approval of that support by Dennett himself.

Dennett (1993:222) cites an excerpt from Millikan, a comment similar to this one:

"It is possible to define the content of an intentional icon with considerable determinacy if you do it in the following way. Consider the content to be that mapped feature to which the icon specifically adapts the user(s) of that icon. It is the feature that if removed from the environment or incorrectly mapped, will guarantee failure for its users"

(Millikan 1993:109)

He goes on to say:

"This supports and illustrates my fundamental point about constraints on interpretation. We are bound to ascribe content that makes it come out that there is a free-floating rationale in support of this design. When determining the semantic value of something, we should refer... to the general principles in accordance to which it was designed to be guided"

Dennett (1993:223)

Now it is a small step to make from there to the ascription of tacit beliefs to systems on the basis of the environmental assumptions they make, and that with or without the presence of a triggering indicator which adapts them to the presence or absence of the crucial conditions. And in the interest of accounting for cases of tacit representation, it is a step I think Dennett should take. 4

Note finally that this talk of environmental assumptions provides a way of answering embarrassing intentional attributions of beliefs and desires to planets in order to explain their behaviour, notably representation of Kepler's Laws and the belief that they are true. Fodor deals with this by saying that: (a) beliefs are explicit representations, i.e. states with a certain content which enter in the causation of behaviour, (b) states representing Kepler's Laws do not enter in the aetiology of the behaviour of the planets, therefore (c) planets do not know Kepler's Laws, believe them to be true or whatever (Fodor 1987: footnote 9 p156). A charge that Fodor might press against Dennett is one to the effect that his liberal position vis a vis representation would allow the planets to tacitly represent Kepler's laws. A response here might be to say that the behaviour of the planets isn't adapted in any sense. The truth of Kepler's Laws isn't then an assumption the planets make when carrying out some proper function, as there is no such function in sight. Another move might have been to add a condition to tacit representation to the effect that it can only be procedural knowledge defined over explicit representations. This is what Dretske (1988) does when he discusses a belief - desire gloss on the behaviour of a rat (O), conditioned to press a lever (M) upon the flashing of a red light (F) (and perception / belief of the event (B)) when hungry (D) to bring about the releasing of some food into the cage (R). Dretske wants to say that the rat does M because he wants R (c.f. the explicit internal indicator D) believes that F is the case (c.f. the explicit internal indicator B) and believes that doing M upon F will bring about R. The problem is that, while there is internal causally potent tokening for the belief that F and the desire that R, there is no such thing as an internal token for the last belief. Dretske resorts to tacit belief, and adds a definition-over-already-intentionally-characterised-objects condition to avoid excessive liberalism (Dretske 1988:117-118). I do not think however that Dennett would want such a stringent condition on tacit representation. Also, his suggestions concerning tacit desires go against that grain of this5.

I am referring here to Dennett's oft-cited report concerning a programmer saying of a chess program "it thinks it should get its queen out early" (Dennett 1979:107). The point is that the instruction appears to guide the chess program's behaviour without any kind of representation involving earliness and getting queens out being tokened. Now both Fodor (1987) and Cummins (1986) have interpreted this as a piece of procedural knowledge. I think it is probably more perspicuously analysed as a tacit desire. The program wants its queen to be out early on in the game. Cummins suggests a control-implicit way of achieving this outcome without explicit tokening. It goes something like: pieces are moved out in increasing order of value (low-ranking pieces preferred) and according to a measure of aggressiveness of the move (aggressive moves preferred), aggressive moves are difficult to make with low-ranking pieces early on in the game, high-ranking pieces thus tend to get deployed early. The idea is that the chess program was designed in such a way as to bring about early queen deployment. Having the queen out early is a function of CHESS's behaviour. The behaviour is explained by reference to a desire, the object of which is the fulfilment of the proper function of the behaviour.

Ascription of tacit desire on the basis of biological interest is an extremely common and intuitive move. Indeed, we do this for a lot of simple tropistic behaviour in which really one could only pick out a single internal state, i.e. the indicator that the crucial conditions for the triggering of the behaviour, in the aetiology of the behaviour. No internal desire-like token in sight. Plant behaviour is another case in point. The rubbing of a roundworm against the inner surface of some predacious fungi triggers off, via some internal indicator, the production of a small loop which garrottes the poor worm (Dretske 1988). Why does the fungus garrotte the worm? Because it wants to kill it, eat it, etc.67

We can now offer a definition of tacit desire on par with the definition offered above in (A):

(B) An organism O which:

(i) has a trait T which has a function to bring about E by means of doing R, and this relying on circumstances A1,...Ai to perform satisfactorily, and

(ii) does not have a trait (a) the function of which is to adapt T by eliciting T (and therefore R) in the sole presence of circumstances A1,...AI8, and (b) which is in a state such that it itself has the function of eliciting T (and therefore R).

is said to tacitly desire that E obtains. It is said to explicitly that E obtains in case (ii) fails.

At this point one might start to see the inkling of a response to (my) Fodor's 'No Intentional Explanation without Intentional Causation', and a clearer idea of what Dennett might have been trying to say when claiming that belief-desire attribution was a matter of laying bare the rationale for behaviour rather than picking out salient parts in its local aetiology.

It is worth here quoting Dennett in full concerning his views on rationality. His closing contribution to Dahlbom's Dennett and His Critics contains an explicit statement which we can only welcome given the lack of clarity in his earlier treatments of this supporting pillar of his work. Furthermore, there is also a much appreciated (by myself) shift away from his early handwaving towards formal theories of inference processes, towards a more function-based approach.

"Millikan's main point of disagreement with my view - or with what she takes to be my view - concerns the role I have given rationality. When she asks if, as my early slogan has it, rationality is the mother of intention, she has in mind a reading of rationality as logicity, starting down the same path as Mc Ginn (and a host of others I am sure). I think this obscures to her the continuum of process-types that take us from brute non-inferential tropism, through various inference-like steps, to full-fledged 'rational thinking'. One can start with the rationality of the design process that produces the mechanisms that do the work. If the thing is wired up rationally, it doesn't need to do any further thinking, it will do the rational

1 thing. It doesn't have to (a) identify a state that has meaning, (b) figure out using "rational thought" what it means, (c) figure out using more "rational thought" what the implications are on this occasion of a message with that meaning, (d) compose a plan of action rationally justified on the basis of this determination, and finally (e) act on it. If it is rationally wired up, the mechanism will take care of all this. That is the sense in which I view rationality as the mother of intention,..."

Dennett (1993:225)

We have a view according to which rational behaviour is adaptive, or adapted, behaviour then. Giving reasons a way of showing how that behaviour is functional. As Dretske puts it, we are explaining why rather than explaining how. In his more idiosyncratic way of putting things, we are providing a 'structuring' rather than a 'triggering' cause (Dretske 1988). Laying bare the function and normal conditions for functioning provides the required environmental embedding.

This idea that rational explanation is explanation takes adapted behaviour as an explanation provides us at least with some kind of minimum cutoff point at which we can say "now we are definitely leaving our domain of explanation, beyond this point, belief-desire attributions can really be no more than illegitimate anthropomorphisms or zoomorphisms, at best we are cooking up spurious attributions to provide a pseudo-explanation or pseudo-prediction to fit our existing knowledge of correlations between observables, we have no independent grounds to make these assertions". This cutoff point, I would suggest, following etiological accounts of function, corresponds to the moment behaviour is selected for, promoting the reproduction of the behavior. From that moment the behaviour acquires a function, and subsequent generations of likewise behaviorers reasons for behaving likewise. It is at this level where, doing some distinction between simply doing as reason dictates (by chance) and doing because reason dictates, can start doing some work (Dennett 1984: 25-27). So exit from intentional systems awkward candidates like electricity, volcanoes and the wind. 10

If the tentative appeal made to tacit belief and tacit desire as relational properties between an organism and states of affairs, grounded in biological function, is anything to go by, we have a Realist position (with definite identity conditions and everything) with which to support an Intentional Explanation without Intentional Causation. Insofar as giving "an account of the 'mysterious connection' between reasons and actions" (Davidson 1963) is concerned, I hope to have convincingly suggested that the causal theory need not be the only naturalistically kosher game in town.

Perhaps then, for Dennett, Bechtel's invitation to Realism through function never seemed more appropriate. 11

3. Synopsis and concluding comments.

We began the paper by examining Dennett's comments to the effect that propositional attitudes can be legitimately ascribed without a corresponding tokening of an 'internal representation'. We then attempted to give a systematic account of the phenomenon by appealing to biological teleology. This left us with an account of what a tacit belief or desire is, namely a relational property between an organism and certain states of affairs which holds by virtue of that organism possessing a trait with a certain function (c.f. (A) and (B) above). But what should we say of explicit beliefs; what is an explicit belief? So far, our account could very well have endorsed some orthodox Fodorian view which might go something like this: having an explicit belief is a matter of bearing a certain relationship to an internal state, a 'belief', which in turn bears a certain relationship to a state of affairs. Notice how this leaves us with an inelegant disjunctive and ontologically heterogeneous account of what a belief is: on one hand a concrete particular, on the other hand a relational property. This is suspect to say the least: what are the features in common to both tacit and explicit beliefs which both make them instances of beliefs? Some tidying up is called for. What I propose is that believing is only a matter of bearing a certain relational property vis-a-vis possible states of affairs. Consequently a belief is not taken to be a concrete particular. Beliefs are never internal states, and are never causes of behaviour. Now of course one might respond "well, if you consider an internal state A, and can legitimately say that the organism who has A as this internal state A believes that p when in state A and doesn't believe that p when not in state A, surely you can't avoid saying that that state is the belief that p?". Well consider: O tends to sink in water when its pouch (say) is full of lead. O tends not to sink in water when its pouch is not full of lead. The pouch being full of lead isn't for that much the tendency to sink in water, is it? My contention is then, at the end of this discussion, that what people like Fodor have termed beliefs and desires are simply features of the organism by virtue of which the organism bears a certain relational property vis-a-vis a certain state of affairs.

Acknowledgements.

I would like to thank Ron Chrisley for helpful comments on an earlier draft of this paper.

References:

- Bechtel, W. 1985: Realism and the Intentional Stance. *Cognitive Science* 9: 473 - 497.
- Brooks, R.A. 1991: Intelligence without Representation. *AI* 47: 139 - 159.
- Cummins, R 1986: Inexplicit Information. In M. Brand and R.M. Harnish (eds.) *The Representation of Knowledge and Belief*. University of Arizona Press.
- Davidson, D. 1963: Actions, Reasons and Causes. *The Journal of Philosophy*. Reprinted in D. Davidson (1980) *Essays on Actions and Events*. Oxford: Oxford University Press.
- Dennett, D 1979a: *Brainstorms*. Brighton: Harvester Press.
- Dennett, D. 1979b: True Believers: The Intentional Strategy and Why It Works. In A.F. Heath (ed) *Scientific Explanation*. Oxford: Oxford University Press.
- Dennett, D. 1981: Three Kinds of Intentional Psychology. In R. Healy (ed.) *Reduction, Time, and Reality*. Cambridge UK: Cambridge University Press.
- Dennett, D. 1982: Styles of Mental Representation. *Proceedings of the Aristotelian Society* vol. LXXXIII.
- Dennett, D. 1983: Author's Response. In *Open Peer Commentary on Dennett's Intentional Systems in Cognitive Ethology*. BBS 3: 369 - 370.
- Dennett, D. 1984: *Elbow Room*. Camb. Mass.: MIT Press.
- Dennett, D. 1993: Back from the Drawing Board. In B Dahlbom (ed.) *Dennett and his Critics*. Oxford: Basil Blackwell.
- Dretske, F. 1988: *Explaining Behaviour: reasons in a world of causes*. Camb. Mass.: MIT Press.
- Fodor, J. 1987: *Psychosemantics*. Camb. Mass: MIT Press.
- Hadley, R. 1995: The Explicit - Implicit Distinction. *Minds and Machines* 5: 219-242.
- Kirsch, D. 1991: When is Information Explicitly Represented? In P. Hanson (ed.) *Information, Thought and Content*. UBC Press.
- McFarland, D. 1983: Intentions as Goals. In *Open Peer commentary on Dennett's Intentional Systems in Cognitive Ethology*. BBS 3: 369 - 370.
- McFarland, D. 1989: Goals, No Goals and Own Goals. In A. Montefiore and D.Noble (eds.) *Goals, No Goals, and Own Goals*. Unwin Hyman, 39-57.
- McFarland, D. 1993: Rational Behaviour of Animals and Machines. *Proceeding of a conference entitled "Prerational Intelligence" Nov 22 - 26 1993, University of Bielefeld*.
- McFarland, D. and T. Bossler 1993: *Intelligent Behaviour in Animals and Robots*. Camb. Mass.: MIT Press.
- Millikan, R.G. 1984: *Language, Thought, and other Biological Categories*. Camb. Mass.: MIT Press.
- Millikan, R. 1993: *White Queen Psychology and Other Essays for Alice*. Camb. Mass.: MIT Press.
- Sehon, S. 1994: Teleology and the Nature of Mental States. *American Philosophical Quarterly* vol. 31(1): 63-72.
- Smolensky, P 1988: On the Proper Treatment of Connectionism. *BBS* 2, 1-74.

1 There is much I disagree with in David Kirsch's article, though given that, as we shall now see, a 'Kirschian' objection to Fodor's claims isn't really going to get us that far and can be easily accommodated, I shall not attempt to give a thorough critique here. Hadley (1995) does a good job of that.

2 A brief comment here on my use of the concept of function. Although I am at present sympathetic to current orthodoxy in terms of accounts of function, namely some version of the so-called 'aetiological account', I wish to remain uncommitted at this point to any particular view on the question - the jury may still be out.

3 I have in mind what would correspond to Millikan's 'indicative intentional icon' producer (Millikan 1984: 96-102) or Dretske's representation producers in 'type III' representational systems (Dretske 1988: 62-64).

4 At this point, one might want to note how the treatment of tacit representation given above differs from that given in Hadley (1995). Hadley includes cases involving environmental assumptions, such as the so-called 'smoothness' assumptions allegedly made by visual systems, under the heading of "broadly defined implicit representation":

"Representation A is implicit in a representation or system B iff A is not explicitly represented by B, and A is derivable from B by means of logico-mathematical inference conjoined with true descriptions of B's structure, and accurate principles of science. The principles of science just alluded to may include biological, perceptual and other principles of interaction between agent and environment."

(Hadley 1995:236)

This seems to me to be erring too much on the liberal side. According to this definition, tire marks on a roadkill might tacitly represent a Michelin 145 SR tire, as the representation is derivable from a true description of the roadkill's structure conjoined with accurate principles of science. Hadley is not capturing here what he presumably set out to capture: the idea of "success conditions" (p235) or environmental features a system is "well adapted" to (p235). 'Derivability' with "accurate principles of science" is just too lax a requirement.

5 We have until now held the view that the organism tacitly believes that the Normal conditions for proper functioning of the behaviour it is carrying out obtain. Beliefs like 'M brings about R', or 'M brings about R in conditions F' might not prima facie fit the mould that smoothly. Is Ms bringing about R a Normal condition for proper functioning of M? M's bringing about R is conditional upon M's bringing about R' does, at first sight, appear to be tautological. Pending a closer examination of the issue, we can settle with the following amended definition: an organism is said to tacitly believe (a) that the Normal conditions for proper functioning of the behaviour it is carrying out obtain, (b) that the behaviour it is carrying out will perform its proper function, i.e. perform Normally.

6 Scott Sehon (1994) defends such a teleological view of desire, according to which an agent is said to 'desire' the telic outcome of its behaviour, offering an alternative to Davidsonian (incl. 'Fodorian') accounts of action. An account of how one is to treat belief within this framework is lacking however.

7 The idea of tacitly represented goals is now monnaie courante in the ethological and situated robotics literature. Brooks' subsumption architecture is a case in point (e.g. Brooks 1991), as is McFarland's work (see e.g. McFarland 1989, but especially McFarland and Bossert 1993: 183; 184-187, in which he and his colleague air their hostility towards Davidsonian action theory, a quick dip into the philosophical debate which is rare in the largely empirically-minded robotics community).

'Dennettian' as he may be, McFarland gets cold feet however when it comes to calling the tacitly represented goals 'desires'. He maintains that 'an intentional system, in essence, contains a representation of a goal (or want) that is in some way instrumental in controlling the behaviour of the animal' (McFarland 1983). In response to this, Dennett urges him to embrace a more liberal view of intentional idioms so as to include both active (i.e. explicit goal-representing) and passive control systems (Dennett 1983: 381).

8 I have in mind what would correspond to Millikan's 'imperative intentional icon' producer (Millikan 1984: 96-102).

9 See also here the views of McFarland who has been recently pushing a view according to which rational action need not be the result of rational thought (McFarland 1993, McFarland and Bossert 1993).

10 Concerning the behaving as versus behaving because reason dictates distinction, note that an aetiological teleological theory has the resources to meet the Davidsonian challenge to distinguish between the reasons which (happen to) justify an action and the precise reasons for which an action was performed (c.f. Davidson 1963:9). Behaviour may incidentally fulfil many functions, but it is the proper function, whose fulfilling by the behaviour in the past accounts for the presence of the behaviour in the present.

11 Bechtel 1985.