

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

On The Possibility of an Extended, Yet Reductionist Mind

Permalink

<https://escholarship.org/uc/item/3k95f1nf>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Author

Ayala, Saray

Publication Date

2013

Peer reviewed

On The Possibility of an Extended, Yet Reductionist Mind

Saray Ayala (saray@mit.edu)

Department of Humanities: Philosophy, Language & Literature
University Carlos III of Madrid - Madrid 126, 28903 – Getafe. Madrid. Spain
Department of Linguistics & Philosophy, Massachusetts Institute of Technology
77 Massachusetts Avenue, 32-d808, Cambridge, MA. 02139-4307. USA

Abstract

In the debate around the extended mind, the special alliance that the extended thesis often has with functionalism usually plays in favor of the former, with functionalism providing support for the extended thesis. Here I want to consider this alliance in the opposite direction: does the extended thesis provide support for functionalism by promoting the need of a level of explanation that is independent of implementational (in particular neural) details? In spite of a seemingly promising line of reasoning for an affirmative answer, I show here that a commitment to the extended thesis or any version of externalism neither paves the way for a functionalist (or any other anti-reductionist) position nor is incompatible with an explanatory reductionism about the mind. I arrive to that conclusion after analyzing an argument by van Eck et al. (2006) meant to conclude the opposite, and showing why it is unsound.

Keywords: externalism; reductionism; functionalism; Marr's computational level.

Extended, therefore abstracted away from implementational details?

When the thesis that the mind extends beyond the limits of the head (in all its different versions: extended mind, extended cognition, environmentalism, wide computationalism, etc. Extended thesis for short) is developed together with some version of functionalism, often the latter provides support for the former, by guaranteeing the legitimacy of a coarse-grain level of analysis where the implementational particularities do not matter (e.g., Clark & Chalmers, 1998; Clark, 2010). According to this extended functionalism, the details of implementation (be it inside of outside the head) are not what decide the mental status of a structure or process, but rather the causal role it plays in the total economy of the system. Thus, the functionalist stance paves the way for an extended mind. Sometimes, though, this union becomes a reason for criticism of the Extended thesis. One strong line of criticism against some versions of the Extended thesis that draw on this alliance, claims that there is no such substantive theory about the (extended) location of minds, but a mere consequence of functionalism (see Wheeler, 2010), which, in fact, does not care about where the implementational base is located. Thus, it seems that the

Extended thesis would be better off by elaborating an argument that does not rely on functionalist justifications. In line with this predicament, a second wave of the Extended thesis has been developed in the last years (Sutton, 2010; Wilson, 2010).

Instead of the beneficial or not so beneficial role that functionalism can play for extended theorizing, here I address the opposite possibility. My goal is to test the potential for externalist considerations in general to legitimize the autonomy of a level of explanation above implementational details, and therefore, to provide support to anti-reductionist views like functionalism. My motivation is the apparent appeal of a line of reasoning that takes externalist considerations to justify the autonomy of such a level, and so it justifies a functionalist position where what matters for mentality is not the physical particularities of a state but its functional profile, usually described in computational terms.

That apparently appealing reasoning goes as follows: “if environmental structures play a critical role in mental processes, then our explanations of them cannot be limited to neural descriptions, since we need to account for those external elements. Therefore in order to explain cognition we need to approach it from a higher level of explanation that abstracts away from implementational details. A perfect candidate for this is a computational level.”

This argument says that the inclusion of external elements shows the necessity of a coarse-grain level of analysis. And so it can be used to argue for an anti-reductionist position. In Clark's version of the extended thesis (Clark & Chalmers, 1998), this coarse-grain level responds to the need to account for what makes internal and external elements mental, which is their computational role in the total functional economy of the system. Here I want to show why this apparently appealing reasoning is wrong. My conclusion is that neither the Extended thesis nor in general any version of (vehicle or content) externalism provide support for the autonomy and legitimacy of an independent-of-implementational-details explanation of mental phenomena. I will focus on a particular case where the above argument has been applied, that is the work of van Eck, Looren de Jong & Schouten (2006).

The environment in visual perception

van Eck *et al.* (2006) provide us with an argument to dismiss explanatory reductionism on the basis of empirical research in visual perception suggesting the significance of the environment in vision, and in general in cognitive processes. Their conclusion is that the inclusion of the environment in cognitive process shows the irreducible nature of psychological explanations operating at the computational level, since only at that level we can account for the environmental elements.

Some relatively recent studies give support to the idea that vision is not exclusively an internal process (e.g., Ballard *et al.*, 1997). The phenomena of *change blindness* (Rensink *et al.* 2000; Simons, 2000) and *inattentional amnesia* (Wolfe, 1999) show, *contra* what traditional representationalist theories claimed (Marr, 1982), that subjects do not build a complete and detailed representation of the outside world. Instead, they rely on the stability of the environment, exploiting its resources. These recent findings could be said to be sympathetic with Gibson's idea that the ecological context is a central component in visual perception (Gibson, 1972). Although representations are still on the arena and Marr's (1982) classical computational theory of vision is not challenged yet, these findings point towards a more direct relation of the perceiver with her environment: the subject is not representing the whole scene (only part of it), but constantly consulting it, as if the information needed to successfully perceive, was, as Gibson defended, (at least partly) already there in the environment.

Norman (2002) proposes a theory where, apparently, Marr's emphasis on representations and information-processing on the one hand, and Gibson's emphasis on action and environment, can be reconciled. Norman's *Dual Process Approach* exploits the anatomical distinction between two brain pathways serving two different functions. The dorsal pathway is for action, while the ventral pathway is for representation (this approach resembles the perception-action model developed by Milner & Goodale (1995), where vision is said to have two different functions, carried out by two separate brain structures.)

We are interested in how van Eck *et al.* examine Norman's approach and conclude it is incorrect. The findings mentioned above, point towards outside the brain, while Norman's approach restricts the action-related aspect of vision to the functioning of the dorsal pathway within the brain. That is, while Norman's proposal recognises that visual perception also serves to guide motor behaviour (perception is here, as Gibson argued, related to action), it is still internalist, situating the ingredients of perception inside the head. Gibson's direct perception proposal, however, includes the organism's environment, and it is this environmental factor what van Eck *et al.* want to rescue from inside the head.

van Eck *et al.* draw on Norman's mistake, presenting it as a confusion between different levels of explanation. Norman seems to be missing the externalist implications of

recognizing the importance of the environment. In van Eck *et al.* proposed clarification I find, however, two mistakes. On the one hand, they confuse ontological and explanatory issues; on the other hand, they make a controversial identification between Marr's computational level and Gibson's emphasis on the environment.

Internalism vs. Externalism

Norman is wrong in keeping everything inside the head. Ecological perception (Gibson's proposal) is not the creation of action-based representations. It is not (only) about internal mechanisms. In van Eck *et al.*'s opinion, "Norman ignores that environment constitutes an additional level above internal processes" (2006, p. 21). According to them, ecological Gibson-like theories cannot be equated with neurophysiological accounts of the dorsal system. Gibson's view "is much broader than the level of (action) representations, because ecological perception also encompasses the environment" (*ibid.*, p. 21). van Eck *et al.* are here moving the debate towards internalism/externalism. Are mental processes exclusively determined by intra-organism facts, or also by external factors? Gibson's view and the findings mentioned above seem to support some version of an externalist position. Therefore any account of visual perception solely in terms of internal mechanisms is missing something (e.g. this account will be unable to explain why the phenomenon of, for example, *change blindness*, happens). The internalist mistake lies in ignoring that other level, the environment. van Eck *et al.*'s conclusion is that the environment constitutes a different level of explanation, distinct from the internalist one.

This reasoning paves the way for their anti-reductionist conclusion. To account for that extra level we need a discipline that is not confined to lower-level processes; a science that is concerned with mental processes and is not restricted to neural, and therefore, internal processes. That is psychology. Psychology explains mental processes at a higher-level (describing them first at a computational level), where not only neural, but also elements outside the organism can be accounted for. Thus, this goes against any reductionist attempt to explain mental processes with a lower-level science like neuroscience. Thus, findings supporting Gibson's view are said to support an anti-reductionist project where psychology is necessary. Let's summarize van Eck *et al.*'s argument in bullet points:

(i) Environment plays a critical role in (visual) perception; in order to understand this we need to consider organism-environment interactions, and not only what is going on inside the organism.

(ii) To account for organism-environment interactions we need a description at a level wider than a merely internalist one, like Marr's computational level, where the function of (visual) perception is described.

(iii) Psychology provides this kind of description, since it draws on explanations at a level wider than that of neuroscience.

(iv) Psychology is therefore necessary to explain visual perception.

It is now time to consider this argument critically.

Two Confusions

From externalism to anti-reductionism

In my opinion, van Eck *et al.*'s argument fails as an anti-reductionist defence. First, they confuse two different discussions, that is, the internalism vs. externalism debate, with the reductionism vs. anti-reductionism debate. It is true, I believe, that they assume that many practices within neuroscience are internalist (e.g., Bickle, 1998; 2003). So in attacking internalism, their argument can count against internalist reductionists (like Bickle seems to be). But reductionism in general, as a not-necessarily-internalist program, is left untouched.

Once we admit that empirical findings suggest the critical role that environment plays in (visual) perception, we have to account for that environmental factor. This however does not in itself warrant the need for psychology. The need to break the limits of the skull in order to describe the *explanandum* does not warrant the necessity of a higher level of explanation. That is, premise (ii) is not as obvious as van Eck *et al.* seem to present it. We can hold an externalist position and at the same time opt for the lower level of explanation. To tell against reductionism van Eck *et al.* would need, besides empirical findings suggesting that environment is critical in visual perception (and in general in any mental process), a conceptual claim asserting that to account for the organism-environment critical interaction we need to do so in an independent-of-implementation level.

The key assumption in van Eck *et al.*'s argument, premise (ii) is, in my opinion, a *non sequitur*. We can recognise that the environment makes a critical contribution to cognitive processes, and we can even say that the environment is a constitutive part of the system (like some versions of the Extended thesis claim), and it would still be a different thing to claim that this contribution constitutes a different level that needs to be accounted for by an autonomous science that is independent of physical details.

Here it is relevant to bring up the distinction between ontological and explanatory questions. The former seeks the constitution of something, while the latter inquires into a proper explanation of something. The ontological question at issue here is what constitutes vision: processes exclusively inside the organism or also states and features outside the organism? The explanatory question asks for how best to explain vision, or what counts as a (good) explanation of it. The explanatory question that concerns us here asks how we should deal (assuming we accept there is such an ingredient) with the outside-the-organism dimension(s) of vision. A particular response to the ontological question does not (automatically) imply a

specific response to the explanatory question. These two questions have to be answered separately.

Thus, from a claim about the constitution of vision, it is a *non sequitur* to state that as a consequence vision has to be explained in a particular way. Moreover, the (explanatory) demand that an explanation of vision requires reference to the environment still does not entail that a description at Marr's computational level is necessary. In order to assert such an explanatory choice, an additional argument is required. The ontological proof (i.e. that environment is part of visual perception) is not sufficient, and neither is the explanatory demand (i.e. that environment has to be accounted for).

In conclusion, for van Eck *et al.*'s argument to work, in particular, for their premise (ii) to be true, they need to prove that the environmental factor needs to be accounted for in a different, higher level explanation, where the visual task is described in a functional way as an information-processing task, and without mention of the particular (physical) component parts. But they only provide evidence for premise (i), that is, for the fact that environment plays a significant role in vision. They rely on Marr's theory to claim that this critical role that environment plays has to be described at Marr's computational level, where the function of vision is described. It is debatable whether Marr's theory is internalist or externalist. In the former case, van Eck *et al.*'s argument fails, since, if visual processes are (taxonomically) located inside the organism, there is no need for this extra level of explanation, and then their next step towards psychology as the science that accounts for that level, does not follow. If we concede that Marr's theory is externalist, according to what I have said above about the two different types of questions (ontological vs. explanatory), they still have not provided any argument for their explanatory conclusion. It is one thing to maintain that environment is part of our visual processes (ontological assertion), and another, to state that to explain vision we need to do it at a higher level (explanatory claim). The former claim (within the internalist vs. externalist debate) does not imply the latter (which is part of the reductionism vs. anti-reductionism debate).

It is important to notice here that the distinction between ontological and explanatory questions does not map onto the two different debates (internalism/externalism, and reductionism /anti-reductionism). I am not assuming neither that the former debate is exclusively ontological, nor that the latter is solely an explanatory matter. On the one hand, externalist considerations might respond to an explanatory concern, i.e. what is the appropriate unity of analysis. On the other, the debate around reductionism might be presented as an ontological question, e.g. are psychological properties reducible to properties of their implementation basis? The argument I am examining here (i.e., van Eck's *et al.*'s), however, takes the ontological aspect of an externalist position to conclude against an explanatory reductionism. The reason why I find it misleading is not that this argument

drags externalism towards explanatory fields, but that it forces externalism to a particular explanatory position.

A wrong appropriation

van Eck *et al.*'s argument also fails for another reason: they wrongly equate Gibson's emphasis on the environment with Marr's emphasis on an independent-of-the-physical-details description where the cognitive ability to be explained is described as an information-processing task. Let us pay some attention to premise (ii) again.

They claim that the environmental factor has to be explained at a level different from the strictly internal where mechanisms responsible for the task are accounted for. They could then simply rely on Marr's model reputation, as classical cognitive science has done, and claim that a description at Marr's computational level, where the function of the cognitive task is described, is necessary to understand that task. But they go further, and seem to pursue legitimating Marr's computational level. Empirical findings point towards the necessity of including the environment in our understanding of visual perception (what we said to be premise (i)), and this partly supports Gibson's proposal. Interestingly enough, in van Eck *et al.*'s argument these empirical findings are assumed to legitimate Marr's computational level, since it is at Marr's computational level where the interaction with the environment is approached (that is the level of organism-environment interaction, where the task at issue is described independently of the algorithmic and physical details).

Once they claim that we have an "ecological level" to be accounted for (that is, an environmental ingredient that, according to them, calls for a new, higher level of explanation above a purely neural one), Marr's upper level appears as the proper place for that explanation. After all, Gibson's theory is too simplistic, or so a fan of Marr's computational complexity would say, to account for the complexity of organism-environment interaction. And so, although it makes a good point about the importance of environment, Gibson's theory falls short to account for the information-processing complexity that is involved.

Thus, van Eck *et al.* go from empirical findings to the necessity of Marr's computational level, *via* the assimilation of Gibson's ecological level into Marr's three-level model. And this assimilation is what attracts my attention. Gibson's emphasis on the environment does not imply another level above the level of representations, but a completely different anti-representational account of perception. Gibson's ecological level, as van Eck *et al.* label it, is not meant to provide a competing, alternative account of the information-processing mechanisms of vision. Gibson's theory is a completely contrasting explanation of perception where there is no room for internal addition and manipulation of information.

It is this major difference between Gibson's and Marr's theories that invalidates van Eck *et al.*'s assimilation of Marr's computational level (where vision is decomposed as an information-processing task) with Gibson's ecological

level (where vision is considered as an interaction between organism and environment and senses are decomposed into their biological component parts). In my opinion, then, van Eck *et al.*'s are using Gibson's emphasis on the ecological level in an incorrect way. The ecological level has nothing to do with Marr's computational level, so it is not correct to use Gibson's ecological level (and the evidence supporting the importance of environment) to legitimate Marr's computational level and the necessity of psychology.

We see that premise (ii) hides an incorrect assumption. Proving the significant role that environment plays is not the same as legitimating Marr's computational level. Premise (ii) (i.e. to account for environmental contribution we need a description at Marr's computational level, where the function of -visual- perception is described) being incorrect, premise (iii) (i.e. psychology provides this kind of description) has no meaning, because it is unimportant whether psychology provides the kind of descriptions that are required at Marr's computational level. Premise (i), let's remember, is only concerned with the necessity of including the environment in our explanation, not with the necessity of providing a particular kind of explanation. Only premise (i) in van Eck *et al.*'s argument, where empirical findings are reported, reveals as arguably correct. And it is easy to see that empirical support for the importance of environment is not evidence for the necessity of psychology. Premise (i) by itself does not support their anti-reductionist conclusion.

Conclusion

By analyzing van Eck *et al.*'s argument, I tried to show that externalist considerations, and the extended mind thesis in particular (understood as a proposal about the location of the object of study in our explanations of mental processes), neither guarantee nor promote an anti-reductionist methodology and the necessity of a higher level of explanation where mental processes are described at a computational level.

If my considerations above are on the right track, van Eck *et al.*'s anti-reductionist argument is a poor one. On the one hand, they are mixing explanatory and ontological questions, extracting an explanatory conclusion from an ontological claim. On the other hand, they wrongly equate Gibson's emphasis on the environment with Marr's emphasis on an independent-of-the-physical-details description where the cognitive ability to be explained is described as an information-processing task. The identification of the first mistake, so I have argued, reveals that they only provide evidence for the ontological claim, leaving the explanatory assertion without any defence. The second mistake draws on the former. The empirical findings suggesting that environment plays a decisive role in cognitive processes supports Gibson's emphasis on the environment, but does not necessarily legitimates Marr's computational level of explanation, where organism-

environment interaction is accounted for. By equating Gibson's emphasis on the environment with Marr's computational level, they are trying to provide the functionalist sympathy for a description at Marr's upper level with empirical support. If their strategy were right, empirical findings pointing towards the significance of the environment would give support to the (computational) functionalist claim that visual perception, and mental processes in general, need to be approached as a platform-free information-processing task.

The empirical findings they mention might or might not be said to be evidence for Gibson's theory. What I think is clear is that they do not give any support to the claim that cognitive abilities need to be described at Marr's computational level. Their anti-reductionist conclusion is not well supported by their argument. It would require more and different argumentation to go from ontological claims about the location of cognitive processes to an explanatory anti-reductionism (e.g. some claim linking the inclusion of environment in explanations of mental processes to the need of accounting for it in an independent-of-the-implementation-details way).

An extended mind, and in general an externalist mind where environmental factors are a key ingredient, does not necessarily call for a high-level of explanation *à la* Marr's computational level. In our quest for cognition, we can go out of the head and still expect to satisfy our explanatory demands at the implementational level.

Acknowledgements

This work has been developed with the support of the Alliance 4 Universities Program and the Department of Humanities: Philosophy, Language and Literature at the University Carlos III of Madrid. Special thanks to TECNOCOG Research Group at Universitat Autònoma de Barcelona.

References

- Ballard, D., Hayhoe, M., Pook & Rao, R.P. (1997). Deictic Codes for the Embodiment of Cognition, *Behavioral and Brain Sciences*, 20, 4: 723-767.
- Bickle, J. (1998). *Psychoneural Reduction: The New Wave*. Cambridge, MA: MIT Press.
- Bickle, J. (2003). *Philosophy of Neuroscience: A Ruthlessly Reductive Approach*. Dordrecht, Netherlands: Kluwer Academic Press.
- Clark, A. (2010). Memento's revenge: The extended mind extended. In R. Menary (Ed.), *The extended mind* (pp. 43–66). Cambridge: MIT Press.
- Clark, A. and Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7-19.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin: Boston.
- Marr, D. (1982). *Vision*, San Francisco: W.H. Freeman.

- Milner, A. D. & Goodale, M. A. (1995) *The Visual Brain in Action*, Oxford: Oxford University Press.
- Norman, J. (2002). "Two Visual Systems and Two Theories of Perception: An Attempt to Reconcile the Constructivist and the Ecological Approaches", *Behavioral and Brain Sciences*, 25: 73-144.
- Rensink, R. A., O'Regan, J. K. & Clark, J. J. (2000). "On the Failure to Detect Changes in Scenes Across Brief Interceptions", *Visual Cognition*, 7: 127-46.
- Simons, D. J. (2000). "Current Approaches to Change Blindness", *Visual Cognition*, 7:1-15.
- Sutton, J. (2010). Exograms and interdisciplinary: history, the extended mind and the civilizing process. In R. Menary (Ed.), *The extended mind* (pp.33-81). Cambridge: MIT Press.
- van Eck, D., Looren de Jong, H. & Schouten, M. (2006). "Evaluating New Wave Reductionism: The case of Vision", *British Journal for Philosophy of Science* 57: 167-196.
- Wilson, R. (2010). Meaning making and the mind of the externalist. In R. Menary (Ed.), *The extended mind* (pp. 167–188). Cambridge: MIT Press.
- Wolfe, J. M. (1999) "Inattentional Amnesia", in V. Coltheart (ed.), *Fleeting memories*, Cambridge, MA:MIT Press, pp. 71-94.