Thrown Machines: Fixing Heideggerian AI

Michael Wheeler
Department of Philosophy
University of Stirling

Meet the Frame Problem

- “The significance to be given to each logical element [each internally represented piece of data] depends on other logical elements, so that in order to be recognized as forming patterns and ultimately forming objects and meaningful utterances each input must be related to other inputs by rules. But the elements are subject to several interpretations according to different rules and which rule to apply depends on the context. For a computer, however, the context itself can only be recognized according to a rule...

  ...[T]o pick out two dots in a picture as eyes one must have already recognized the context as a face. To recognize this context as a face one must have distinguished *its* relevant features such as shape and hair from the shadows and highlights, and these, in turn, can be picked out as relevant only in a broader context, for example, a domestic situation in which the program can expect to find faces. This context too will have to be recognized by its relevant features, as social rather than, say, meteorological, so that the program selects as significant the people rather than the clouds. But if each context can be recognized only in terms of features selected as relevant and interpreted in terms of a broader context, the AI worker is faced with a regress of contexts.” (Dreyfus, *What Computers Still Can’t Do*)

The Frame Problem continued...

- The question: how is an intelligent agent, conceived as a purely mechanistic system, able to home in on just those aspects of all the things it senses, knows or believes that are relevant in the present context of activity, while ignoring everything that is contextually irrelevant, and how is that agent then able to revise or act on that information in a contextually appropriate manner?
- Dreyfus writes: “Wheeler rightly thinks that the simplest test of the viability of any proposed AI program is whether it can solve the frame problem.” (All further quotes from Dreyfus are from *Heideggerian AI*)

Delaying the Inevitable

- One immediately attractive response is to claim that intelligent systems should deploy stored heuristics or representations of context that determine which of their stored bodies of information are relevant in the present situation.
- But how does the system decide which of its stored heuristics or potentially context-specifying representations are relevant? Another, higher-order set of heuristics or representations would seem to be required, and the same problem seems to re-emerge.

The Frame Problem in Two Dimensions

- Intra-context frame problem: challenges us to say how a purely mechanistic system might achieve appropriate, flexible and fluid action within a context.
- Inter-context frame problem: challenges us to say how a purely mechanistic system might achieve appropriate, flexible and fluid action in worlds in which adaptation to new contexts is open-ended and in which the number of potential contexts is indeterminate.
- So how might the frame problem (both of them!) be addressed?

Cricket Lessons

- Webb's account of cricket phonotaxis.
- There is "no need [for the female cricket] to process sounds in general, provided [the male cricket's signal] has the right motor effects. Indeed, it may be advantageous to have such specificity built in, because it implicitly provides 'recognition' of the correct signal through the failure of the system with any other signal" (*Modeling Biological Behaviour*, p.1092)
- Here context is not something that inner mechanisms must reconstruct, once they have been triggered. Rather, context is something that is always there at the point of triggering, woven into the adaptive fabric of the special-purpose extended mechanism.

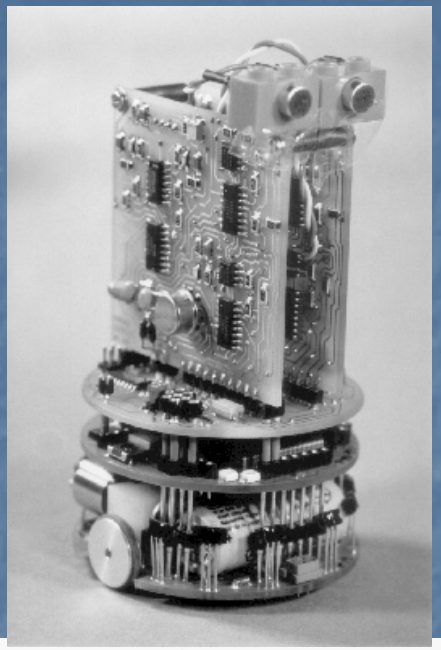

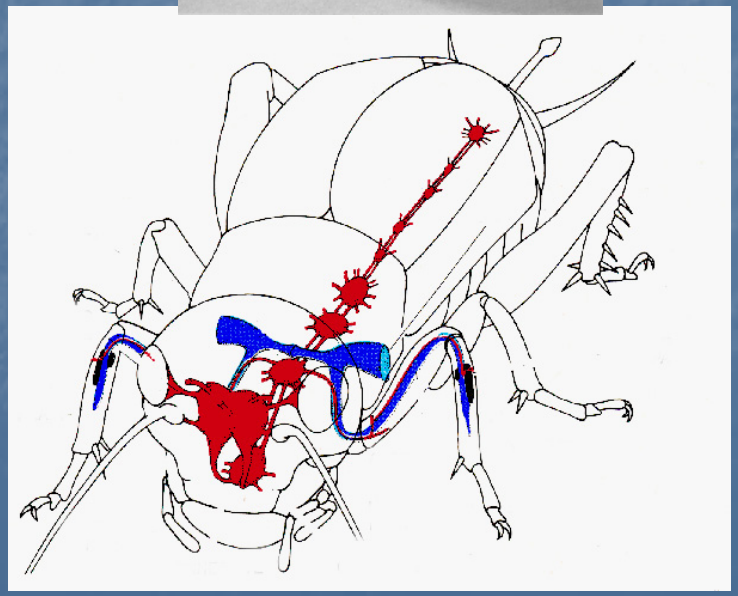

Cognition on the Fly

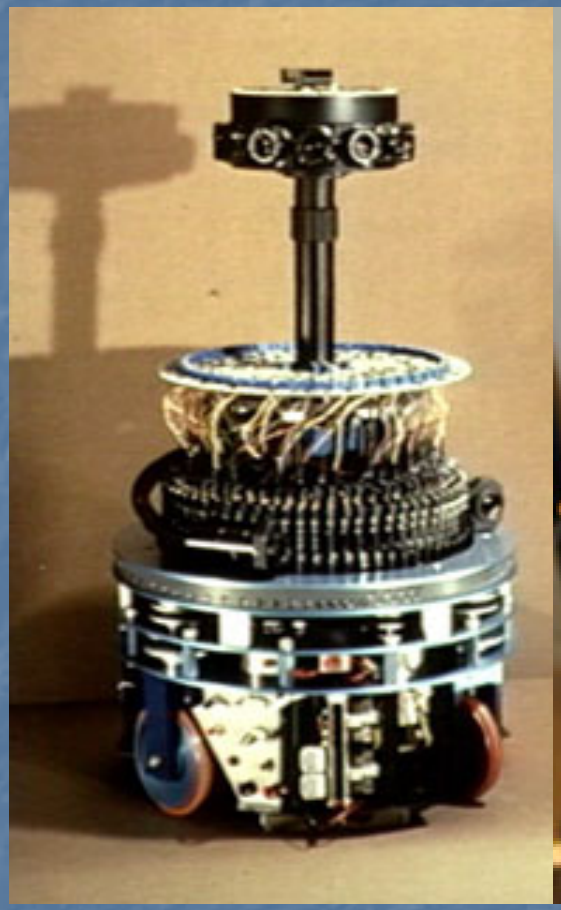

Franceschini, Pichon and Blanes, *From insect vision to robot vision*

Action-Oriented Representation

- The ways in which objects are represented by the fly robot are deeply dependent upon the robot's situatedness and its context-embeddedness.
- The shape, absolute position, and/or orientation of objects are neither calculated nor stored.
- Objects (other than lights) are located by contrast points in the optic flow and represented as avoidance-regions in an egocentrically defined space.
- The representations are action-specific, egocentric, and intrinsically context-dependent (the subagential mark of thrownness).

From a Phenomenological Point of View

- So (I suggest) the intra-context frame problem may be solved by these sorts of situated special-purpose adaptive couplings, perhaps featuring action-oriented representations, perhaps not.
- This is a Heideggerian AI. The subagential profile I have sketched has a direct analogue in agential-level Heideggerian phenomenology
- See Heidegger's account of skillful coping: that's smooth coping in the domain of the ready-to-hand and practical problem solving in the domain of the un-ready-to-hand.
- And his account of thrownness: we are always already embedded in some meaningful context, so we are never in the position of having to add contextual significance to context-independent primitives.
- But we still need to solve the inter-context frame problem

Continuous Reciprocal Causation

- Continuous reciprocal causation (CRC) is causation that involves multiple simultaneous interactions and complex dynamic feedback loops, such that (a) the causal contribution of each systemic component partially determines, and is partially determined by, the causal contributions of large numbers of other systemic components, and, moreover, (b) those contributions may change radically over time.
- See e.g. neural modulation via gas diffusion in the brain as modelled by GasNets (Sussex Evolutionary and Adaptive Systems Research)

A Glimmer of Hope

- The non-representational process of CRC endows a system with a powerful kind of ongoing fluidity, one that involves the functional and even the structural reconfiguration of large networks of components.
- Could this be the mechanistic basis of fast and fluid context-switching?
- CRC and action-oriented representation cannot be present in a single mechanism simultaneously, but nothing rules out their adaptively beneficial co-existence in the same mechanism over time (e.g. transient modularity in some GasNets), or in different mechanisms simultaneously.

Coping Trouble

- Dreyfus accuses me of propagating a “cognitivist misreading of Heidegger” that leads me astray in my handling of the frame problem.
- He argues that we should reject all forms of representationalism about intelligence. As he puts it, “for Heidegger, all representational accounts are part of the problem”. Why is this? Holism, know-how, and intrinsic context-independence
- Here Dreyfus draws a distinction between skillful coping and background coping.
- Background coping is “an even more basic nonrepresentational holistic coping that allows copers to orient themselves in the world”.
- “The important point for Heidegger, but not for Wheeler, is that *all* coping, including unready-to-hand coping, takes place on the background of this basic non-representational, holistic, absorbed, kind of intentionality, which Heidegger calls being-in-the-world.”

…and the Frame Problem

- On the basis of our constantly honed background know-how, we respond directly to relevance, with context-bound entities soliciting or summoning us to act in ways shaped by our past experiences.
- The capacity for flexible context-switching is explained by the fact that I can be summoned not only by the present situation, but also by other situations that, because they have been relevant in the past, lie on the horizon of my experience.
- Thus a "Heideggerian Cognitive Science would require working out an ontology, phenomenology, and brain model, that denies a basic role to any sort of representation" (Dreyfus, *Heideggerian AI*).
- Here Dreyfus appeals to the work of neurodynamicist Walter Freeman

Looking for the Disagreement

- I agree with Dreyfus that background coping is the underlying phenomenological structure on the basis of which fluid and flexible context-switching is possible.
- Still, as Dreyfus himself observes, the dynamical processes that, according to Freeman, are realized in the brain are themselves plausibly interpreted as an instantiation of CRC.
- And I have argued that that CRC is nonrepresentational in character, so to the extent that CRC is the mechanistic basis of background coping, I agree that the cognitive science of that phenomenon will deny any basic role to representation.
- So where exactly is the disagreement?

Wheeler's Problematic Ambivalence

- According to Dreyfus, I remain problematically ambivalent about which model of cognitive mechanism, CRC or AOR, is ontologically more basic.
- Dreyfus thinks this because he thinks that CRC causally underlies all background coping, providing a causal dissolution of the intra-context frame problem as well as the inter-context version.
- And why does he think that? "Wheeler's own proposal... by introducing flexible action-oriented *representations*, like any representational approach, has to face the frame problem head on".
- But Dreyfus overestimates the contribution made by CRC.
- CRC mediates the transitions between mechanisms that individually feature the previously identified property of intrinsic context-dependence.
- Intrinsically context-dependent mechanisms do not face the difficulties of assigning relevance or of representing massively holistic networks of contextual significance.

Wheeler's Not-So Problematic Ambivalence

- So does this mean that I am ambivalent about whether AOR or CRC is ontologically more basic? No!
- It is not action-oriented representation that heads off the intra-context frame problem, but the presence of situated special-purpose adaptive couplings. Representations may figure in such mechanisms, but they may not.
- CRC heads of the inter-context frame problem. So there is a clear sense in which CRC is ontologically more basic.
- If I'm ambivalent about anything here, it's about whether CRC or situated special-purpose adaptive coupling is ontologically more basic, since both are aspects of the structural causal basis of absorbed coping.
- Cf. Heidegger on the world versus the worldhood of the world

Representations: a Mixed Message

- Action-oriented representations inherit context-dependence from the intrinsically context-dependent mechanisms in which they figure, which means that there is no need for context to be represented.
- And to the extent that what is being represented by action-oriented representations remains knowledge *that* the environment is thus and so, it is not doing the frame-problem-inviting job of explaining how we are equipped with the capacity to be flexibly sensitive to what is relevant in a particular context.
- But this way of protecting representational explanation has a profound flip-side, since it means that representations are neither aspects of the background coping that grounds our flexible sensitivity to context-dependent relevance (agential level), nor are they causally explanatory of it (subagential level).