Robots, Action, and the Essential Indexical

Paul Teller
prteller@ucdavis.edu

1. Preamble. Rather than directly addressing Ismael's The Situated Self I will present my own approach to some of the book's issues. Studying The Situated Self inspired much of what follows, and I believe that a close comparison will show that in many respects I am doing no more than instantiating some of Ismael's insights. Where Ismael works at an extremely high level of abstraction (6) I will work as concretely as possible. I believe that I reach some similar conclusions, and while space does not permit detailed documentation, I will make a few observations of close contact. At best what I present will correspond to only a tiny fraction of the many difficult issues that Ismael discusses. All the same, I hope that this little presentation will illustrate how, perhaps, quite a bit of her material might be redone at a lower level of abstraction.[i]

2. Objectives. I will examine two problems. The first, as discussed by Perry (1979), concerns understanding how agents make use of, make applicable to themselves, information presented in third person representations. Perry rubs our noses in the problem with the case of the messy shopper:

> Shopping in a supermarket Perry notices a trail of sugar, evidently coming from a torn sugar bag in someone's cart. He searches the supermarket to find the messy shopper to alert her/him to the problem. What does Perry have to learn to discover that HE is the messy shopper? It would seem that no third person description by itself would do the job. If Perry learns that the man with the loud shirt, the man he sees in the mirror, even the man named "Perry" is making the mess, such information presented in the third person will not, by itself, suffice. Perry must also know that HE is the one in the loud shirt, the one he sees in the mirror, the one named "Perry".

Obviously the problem generalizes. One seems driven to the conclusion that to make information presented in the third person relevant to action, agents must also have something presented in the first person – "I am the one in the loud shirt, seen in the mirror, called Perry". This view has become known as the doctrine of "the essential indexical."

The second problem concerns what is called "demonstrative content" (Peackock, 1981): It is clamed that to form an intention an agent must incorporate information that can only be presented with a demonstration. For example, an agent wants a certain pill bottle but can't organize herself to act merely on the basis of some third person description of the bottle or of its location. In addition the agent must come to appreciate of some perceptually present object that *that* is the pill bottle I want, *that* is the pill bottle satisfying… description, or located at *that* place.

Just what is "demonstrative content"? How does demonstrative content work?

These problems would appear, at least, to be closely related. A third person description of the target of an intention won't do because it does not, by itself, related the agent to the target. Suppose that a description of the pill bottle's location is given in the third person: e.g., its coordinates, xyz, in some fixed coordinate system. Absence of some additional information available in the first person, it is claimed, leaves the agent helpless. The agent must, for example know where S/HE is located in the same coordinate system. Which is just the messy shopper problem. As before, the point quickly generalizes

I will establish that messy shopper and demonstrative content cases do not, by themselves, show that demonstration and use of the first person are ineliminablly required for acting or, if intentions are required for action, for forming intentions.[ii] And thereby I will undermine any use to which some might try to put such claims. To illustrate, consider an argument form that someone might be temped to use to argue for some special status of individuals, dualism, generally what some have called "non-objective", or "perspectival", or "extra-natural" facts. That is, I will have undermined arguments of the form:

1) All "objective" facts can be expressed in the third person.

2) The messy shopper problem and the need for demonstrative content involve information about facts that cannot be expressed in the third person.

3) Therefore there must be some kinds of "non-objective" facts.

I will establish that any such argument is unsound by providing toy models that will illustrate exactly what premise (2) claims can't happen. There is no reason to think that people operate in any way like the models that I will describe. But since the contentions arguments are "can't possibly" arguments, "how possibly" examples suffice to undermine them wholesale.

3. Methodology. I will systematically state everything in the third person and without indexicals or demonstratives of any kind. This will sometimes require complicating the exposition and using awkward circumlocutions. But in this way I insure that "first person" and "demonstrative" information has not been smuggled in.

I will be describing robots, using agentive terminology: agent, act, choose, deliberate, perceive… Some will claim that agentive terminology for robots is metaphorical only, that robots no more act than do mouse traps and thermostats, so that the argument will show nothing about what is properly called action by an agent.[iii]

I will address this issue at the end. In the meantime, not to beg any questions, I will use the following terminology:

Act (with upper case 'A'): what people do.

Robot-act: what robots do.

act (lower case 'a'): can refer to Actions or robot-actions.

Agent (with upper case 'A), robot-agent, agent, Deliberate, Instruct, etc, similarly.

In extended passages about robots I will take the "robot" qualifier for granted

4. The robots, their environment, and the essential indexical. We consider a collection of robots that move around on a grid with squares marked A1, A2,…, B1, B2,…. The grid has mirrors on all sides, which enable each robot to see all the robots on the grid, including the robot making the observation . Each robot is painted a distinctive color. The components of each robot include a perceptual module that can form third person descriptions and maps of the environment, an action module that executes instructions, and a deliberation module[iv] that has been programmed with an algorithm that takes instructions from the outside and input from the perceptual and other modules, and issues instructions to various modules. Whether or not two modules are both modules of the same robot is fixed by whether the two modules are wired to each other.

Instructions to the robots periodically appear on a message board. Red-robot: move to B5. Green-robot: move to F11. Etc. Each robot's perceptual module periodically scans the message board and forwards any new instruction to the deliberation module to which the perceptual module is wired. To avoid conflicting instructions, we suppose that only one instruction at a time is displayed on the message board, and when the instruction has been followed, the instruction is erased. This also insures that only one robot moves at a time. When a deliberation module receives an instruction, the deliberation module executes a routine that determines whether the color of the robot that has just received the instruction is the same as the color of the robot addressed in the instruction. Two different procedures provide two different versions of the set up.

Method one: personal names or descriptions (Price, 1989 p. 242). Each robot has a distinguished register called the "identification register". The red robot's identification register contains "red", the green robot's "green", etc. Each deliberation module determines whether or not to perform the robot-action described in an instruction by retrieving the color-name in the identification register: If the color-

name in the register agrees with the name in the instruction, the deliberation module performs the instruction. Otherwise not.

Method two: Even personal names or descriptions are not required. The deliberation module, on receiving an instruction, issues an instruction to the action module: Stamp feet! and instructs the perceptual model to note and report the color of the robot that has just stamped its feet.(Recall that only one robot moves at a time.) If the color reported is the same as the color mentioned in the instruction, the deliberation module executes the instruction. Otherwise not.

Using no more than third person representations each robot executes an instruction just in case the robot that has received the instruction is the robot to which the instruction is addressed. There has been no use of anything like "I am the red-robot", "This perceptual module is the one in the red-robot", or "The red-robot has determined that it is the one receiving the instruction." For robots use of indexicals is not essential! The example does not show that the way in which real Agents – in particular humans - perform Actions does not require use of indexicals, should robot-action and Action not coincide. I will return to this issue.

5. Expanding the use of information represented in the third person: Solving the messy shopper and demonstrative content problems for robots. So far instructions are the only information presented in the third person to which the robots respond. The case of the messy shopper concerns use of factual information. While I will discus only the one case, I take it to be amply clear that a great deal could be treated in the same way.

Suppose that the message board also periodically posts facts represented in the third person, e.g., "The red robot is making a mess by trailing sugar from a torn sugar bag" (a "torn sugar bag alert"). Each robot's perceptual module picks up the torn sugar bag alert and, still expressed in the third person, passes the alert on to the deliberation module. The deliberation module is programmed to activate the sugar bag patching module (a submodule of the action model just in case the deliberation module determines that the color of the robot of which the deliberation module is a part is identical to the robot-color named (by definite description) in the torn sugar bag alert. And the deliberation module determines the color of the robot of which the decision module is a part by activating the personal-name-lookup or the robot foot-stamping routine.

Messy shopper problem solved. All robots respond appropriately to the third person report that the red robot is making a sugar-mess. No indexical or demonstrative of any kind has been used. Once again, by example I have proved that an indexical is not essential for robot-action.

Turning to a case that would appear to call for "demonstrative content", suppose that the red robot receives the instruction, Red robot: move to the square occupied by the green robot. One might argue that facts represented only in the third person cannot suffice. Using its perceptual module the robot can produce a third person description of the green robot's location, square B5, for example. But to get itself to B5 the robot must first know where IT is. For which, the dogma seems to be, third person description will not suffice. (Where I have here violated the injunction against use of indexicals, to make clear how the argument is alleged to go.)

The red robot chooses, or calculates, the robot-action that will get the red robot to the location of the green robot as follows. First the red robot determines whether or not to perform the instruction as above If there is a positive result the deliberation module instructs the perceptual model to establish a third person description of the green robot's location: "Green robot at F11", and to pass the result back to the deliberation module. The deliberation module now issues two further instructions: to the action module to stamp the robot's feet and to the perceptual model to read the square location of the robot stamping its feet. (Recall that each robot can see all the robots in the mirrors and that only one robot moves at a time.) The perceptual module passes back the result, "Red robot at B5", as it might be, to the deliberation module. Finally the deliberation module now issues the instruction to the action module "Move from B5 to F11." The action module, perhaps also using a third person description of the orientation of the robot with the stamping feet, generates and executes a list of motions that get the red robot from B5 to F11 - move 6 squares forward, then 4 squares to the right.

Thus the robot exhibits the function of "demonstrative content": Given the robot's architecture and present observations, the pair of third person descriptions (B5, F11) serves the function of "THAT robot", that is the robot that is located 6 steps in front and 4 steps to the right of the robot generating the instruction of how to move.

Any sensibly designed robot will be organized in a more flexible and intricate way. I have given stylized examples to make clear that such a task can be accomplished by processing only facts described in the third person.

Note that throughout this account I have leaned on facts about how the robots are wired together and the potential behavioral significances of such structure. In this way my account illustrates Ismael's appeal to "unarticulated constituents" and "architecture" that together do work like that done by indexicals. This shows that, at least in cases such as the above, the kind of things accomplished by using first person statements and demonstratives can be incorporated in the "natural order", as Ismael aims to do (p. 1, 231). And, I suggest, in ways that constitute concrete instances of some of Ismael's much more abstract characterizations.

6. Connection with self-location and self-modeling. In the foregoing the deliberation module that is housed in the red robot has determined the location of the red robot, illustrating the centrality of the idea of self-location, just as Ismael claims. For vividness suppose that instead of location descriptions, A1, A2… the robot uses a map with the coordinates written in each square. The deliberation module draws a red circle around the square with the coordinates that the perceptual model reports as the one

observed with red stamping robot feet. This self-location is accomplished using only third person representation, and "architecture" – facts about how the robot is programmed and put together. Again verifying at least for this special case, that the work usually done with 'I' is accomplished in a completely naturalistic way.

Broadly, the strategy for doing the work of "I" using only third person descriptions has turned on the fact that robot self-location can work just like location of any (other) robot. We describe the way in which any robot establishes the location of any (other) robot and then consider the special case when locator and located coincide.

Similar comments apply for demonstrative content: Again, for vividness the deliberation module may draw a green circle around the square that the green robot is observed to occupy. Together the red and green circles, given the robot's architecture, serve the function of targeting "that robot".

The idea of self-locating generalizes to that of self-modeling. Designers produce much more sophisticated robots by equipping each robot with the power to construct a model of both the robot's environment and the robot's location within that environment. This works with location in physical space, but also with location in other "spaces", for example the spaces of shapes and of weights. Designers achieve greater power yet by generalizing both ideas of modeling and of locating to that of self-modeling of the robots' own internal states.

7. Consolidating the argument. I have shown, by example, that the use of indexicals is not necessary for robot-action. I have not shown that such use is not required for Action. should robot-action and Action not coincide What, then, are the implications of the robot examples for the question of whether real Action requires use of indexicals?

Remember that I claim to establish that messy shopper and demonstrative content type of cases do not, by themselves, show that indexicals are essential for performing Actions. If one has some independent reason for maintaining that robot-action is not Action, only the messy shopper and demonstrative content cases together with this independent reason support the necessity of indexicals for Action. On the other hand, insofar as the coincidence of Action and robot-action is an open question, then whether the robot examples also apply to Action is equally an open question. The original cases, by themselves, do not settle the issue.

As a corollary, if one wants to appeal to the messy shopper and demonstrative content cases as the basis of an argument to show that robot-action is not real Action, such an appeal explicitly begs that question.

Yet another line of argument, not yet broached, resolves the issue of Action vs. robot-action, while sidestepping the whole question of whether Action somehow, on its own, counts as machine-like.[v] [vi] Even supposing that machines cannot, on their own, perform Actions, this supposition will not apply to robots taken as part of a designer-robot system. In such systems, properly designed robots really Act by Acting for their designers by proxy. "Properly designed" here means instantiating a process that we would interpret as deliberation in the way that I will now explain.

Metaphorically, the idea is that intentionality trickles down from designers to their properly designed robots.[vii] I have written as if various aspects of robots bear intentionality very broadly. For example I have written in the spirit that various sentence tokens bear their natural intended interpretations, as for example in: "Red robot: go to F11", and "The red robot is making a mess." Such tokens surely do not bear their natural interpretations in the case of my very simple robots merely by their mechanical operation, and perhaps not for any robots, no matter how skillfully designed. But, even in the simple cases we interpreted these tokens in the natural way, and so in the larger context that includes the designers the tokens bear these natural interpretations. Likewise, one way I have of Acting is to deploy robots – machines - to Act for me, to achieve my ends in ways functionally similar to ways in which I would Act. No matter that I leave the details of which Actions, finally, get performed up to the working of the machine. Patching up a torn sugar bag will count as an Action should I, or my robot, have

the one with the tear, as will reaching for a pill bottle that I, or my robot needs for proper functioning – as I interprete what counts as proper functioning for the robot.

Nothing about the interpretation of robot-motions as real Actions interferes with the way my robot examples work. Hence I have shown that there are Actions – real Actions - that make no use of indexicals.

References

Kobes, Bernard W. (1990) "Individualism and Artificial Intelligence" Philosophical Perspectives, Vol 4. 429-459

Peacock, Christopher (1981) "Demonstrative Thought and Psychological Explanation." Synthese 49, 187-217

Perry, John (1979) 1979. "The Problem of the Essential Indexical." *Noûs*, 13: 3-21

Price, Huw (1989) "Action Explanation and the Nature of Mind." Peter Slezak and W. R Albury (eds), Computer, Brains, and Minds. Kluwer. 221-251

Spencer, David Ian (2009) PhD Dissertation, University of California at Davis

[i] I also have great indebtedness to Huw Price's (1989, see especially pp. 240-242) as well as to work by Ian Spencer (2009). Much of what follows takes Spencer's approach to the problem of using temporal information and transposing it to parallel issues. I was also greatly aided by discussion with Bernard Molyneux and Joshua Earlenbaugh, and by the hard time given me by Bas van Fraassen.

[ii] The examples will bare directly only on action.

[iii] Bas van Fraassen has pressed this worry.

[iv] More carefully called "robot-perceptual, "robot-action" and "robot-deliberation" modules. As mentioned, the "robot" qualifier will be suppressed in extended passages about robots.

[v] Since the whole argument of The Situated Self is to show that intentionality can be treated naturalistically, how "we can integrate our mental lives into the closed causal order described by physics" (231), Ismael will feel no need to appeal to the following considerations.

[vi] The following argument is inspired by Kobe's (1990)

[vii] I don't think that it is a good objection to suggest that the use of indexicals trickles down with the intentionality. At the very least, there would be a problem of question begging similar to one mentioned above.