

## 10 — Reckoning and Judgment

It is finally possible to explain the title of the book.

I have tried to sketch some of what is necessary for a system, human or machine, to be able to think about or be oriented toward the world beyond it—a world with which it copes, through which it navigates, in which it conducts its projects, about which it reasons, to which it is committed, to which it defers. The system must not only be embodied and embedded in this world; it must also recognize it *as* world. No less is required in order for it to distinguish appearance from reality, and choose reality; to recognize the difference between right and wrong—and choose right; to distinguish truth, falsity, and impossibility—and choose truth. No less is required in order for it to register entities, phenomena, people, and situations as such. It must register all these things in the world to which it holds itself, and all that it understands, accountable.

Such, to our inestimable benefit, is the legacy of the human epistemic achievement.

### 10a · Animals

I have not yet said anything about animals.

Nonhuman animals, I take it, are in one sense accountable to the world. There is a huge variety of kinds, needless to say; the forms of accountability pertinent to cockroaches will differ substantially from those applicable to jaguars, or to chimpanzees. Just how much variation there is—how

much difference, how much similarity—I leave to others. And any such account must include the fact that we are animals too. No theory of humans and other animals can ignore the continuity of our constitution, evolution, and development.

How much we overlap with nonhuman animals (especially higher primates) on capacities that contribute to our intelligence is another subject on which I profess no special expertise. It is certainly true that nonhuman animals manifest species-appropriate forms of awareness. Some have exquisitely subtle and precise perceptual (or ‘perceptual’) systems.<sup>1</sup> There are stakes for them; if they get things wrong, they may not survive, individually or as a species. They manifest their own forms of care, and are emotionally sensitive. As is increasingly remarked, in certain respects nonhuman animals outstrip us humans.

Nevertheless, I do not believe that nonhuman animals engage in the sorts of existentially committed practices of understanding of the world *as world* that I have been exploring here. It seems unlikely to the point of impossibility that they are capable of achieving the sorts of objectivity necessary in order to take an object as an object, and thereby to be truly intelligent in a human sense. We may say that a pet recognizes some well-loved object; but that is an object *for us*. Without the four-fold structure of normative commitment laid out above, which a pet seems unlikely to be bound by, it cannot be an object *for them*.<sup>2</sup>

---

1. Presumably backed by massively parallel neural networks.

2. Or maybe pets and creatures can and do register objects as objects, in at least a rudimentary sense. My point is not to defend any particular account of animals, but to get a sense of what is required in order to take there to be objects, and a world—a sense in terms of which we can then have a discussion of whether, and if so to what extent, various types of animal and machine can do that.

Two things can be said for nonhuman animals, though, even if they are incapable of registering an object as an object. First, they *participate* in the world—perhaps even existentially, at least in a rudimentary sense, each in whatever sense that makes sense for them. Second, and crucially, to the extent that they represent or register the world, they *do so in ways that mesh with how they engage, navigate, cope with—and are vulnerable to—it*. The world is presumably intelligible to them, even if not as world, exactly in as much as, and to the extent that, it matches their forms of life, stakes, biological needs, etc.

What matters about animals, that is—and for the present discussion as well—is the (evolutionarily hewn) fit between (i) how they take the world, or whatever patches of it they take, and (ii) how they conduct their lives, what they care about, and what they are vulnerable to. That makes something about their lives or even cognitive powers—if we want to use that phrase for nonhuman animals—*authentic*. I believe that much posthumanist writing overestimates their epistemic powers, but to the extent that animals do something like thinking and understanding, it would certainly seem that they do so in an authentic way.

I will use the term **creature** for systems, including animals, whose registrational powers do not extend beyond the ways in which they conduct their lives, what they care about, what they are vulnerable to, what they are existentially involved with. The term is useful because the development of AI-based “pets” (such as Sony’s Aibo) suggests that we may be, or anyway may soon be, constructing computational creatures in this sense. As I have said a number of times already, it seems not impossible that synthetic creatures based on machine learning, active sensors and effectors, and so on, might not achieve something in the

realm of the authenticity that I have said animals have evolved to possess. That is, nothing in this book argues that it may not be possible to construct genuine computational *creatures* (even if what it would be for constitutive norms to apply to synthetic pets, not merely for them to be at the mercy of our whims, remains obscure).

### 10b · Computers

*The same is not true for computers in general.* Most of the computational systems we construct—including the vast majority of AI systems, from GOFAI to machine learning—represent the world in ways that matter to us, not to them. It is because of that fact that we call them *computers*, or *information processors*, and also because of it that they have power in our lives, that they matter to us. What limits them is that, so far, nothing matters *to them*. To use a phrase of which Haugeland was fond: *they don't give a damn.*<sup>3</sup> Things will only matter to them when they develop committed and deferential existential engagement with the world (Dasein, perhaps, if we wanted to speak that way).

A possibly useful diagram of the difference is given in figure 8—a structure in terms of which to understand humans, nonhuman animals and other creatures, and most of the computers we have built. The forms of life that I have labeled “authentic” lie in the region of creatures, animal or machine, that register the world (or that we register as registering the world) in ways that are appropriate to, *and do not exceed*, the forms of their existential engagement in the world.<sup>4</sup>

---

3. Zed Adams and Jacob Browning, eds., *Giving a Damn: Essays in Dialogue with John Haugeland* (Cambridge, MA: MIT Press, 2016).

4. What is crucial to the creature's engagement is not the registrations themselves, but *the world being as the system registers it*, and the crea-

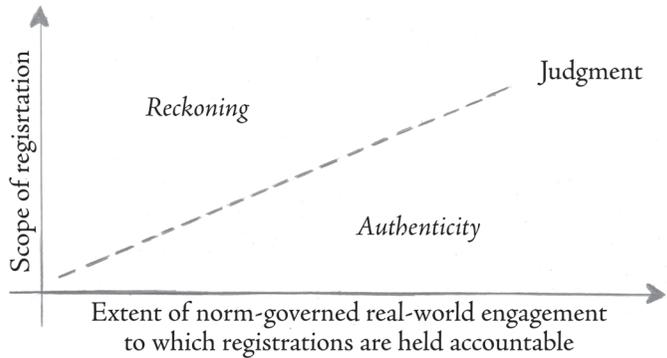


Figure 8

One way to sense the difference is to consider how different are our interactions with nonhuman animals from those with typical AI systems, the outputs of data mining services, and so on. Though we talk to animals, the fact that they do not use high-level language<sup>5</sup> keeps present before our minds the fact that their repertoire of feelings, understandings, projects, and the like, remain constrained to the realms in which they are able to participate authentically.<sup>6</sup> That is exactly not true when our iPhone “tells” us that there has been an accident on the highway, a medical

...  
 ture’s being existentially at stake to that fact.

I would not begin to suggest that this characterization could stand as a definition of authenticity. It is trivial to construct examples of systems that meet the criterion of being “under the line” in figure 12 that no one would count as authentic.

5. At least, so far as we know, they do not use compositional, productive, systematic languages (chapter 7, note 1, pp. 71–72).

6. If we register a nonhuman animal, such as a pet, in ways that do exceed that which is appropriate for how it engages with the world, and do so without intending to be metaphorical or ironic—such as when I say “my dog is frustrated with postindustrial society too”—then we are guilty of attributing a nonauthentic state to it.

system 「claims」 that nausea affects 32.7% of patients who take a given drug, or a wearable monitor 「suggests」 that if you have another drink you should not drive home. The fact that these systems 「use natural language」 can obscure to us the fact that they understand neither traffic nor nausea nor drunk driving. What is essential is that the increasing familiarity of such cases not blind us to the profound difference in types of system.

As noted in the introduction, I use the term “**reckoning**” for the representation manipulation and other forms of intentionally and semantically interpretable behavior carried out by systems that are not themselves capable, in the full-blooded senses that we have been discussing, of understanding what it is that those representations are about—that are not themselves capable of holding the content of their representations to account, that do not authentically engage with the world’s being the way in which their representations represent it as being. Reckoning, that is, is a term for the calculative rationality of which present-day computers (including second-wave AIs) are capable.

I reserve the term “**judgment**,” in contrast, for the sort of understanding I have been talking about—the understanding that is capable of taking objects to be objects, that knows the difference between appearance and reality, that is existentially committed to its own existence and to the integrity of the world as world, that is beholden to objects and bound by them, that defers, and all the rest.<sup>7</sup>

The terminology can be justified in various ways, beyond

---

7. It is because they do not exemplify this suite of properties, as I said in chapter 7 (p. 76), that current computer systems, including those constructed on the principles of second-wave AI, “do not know what they are talking about.”

Hobbes's claim that "reason is nothing but reckoning"—a sentiment that many in AI believe, but with which it is clear that I disagree. But my aim is simple: merely to contrast reckoning with what we convey in plain English by saying that someone handled a situation with "good judgment."<sup>8</sup> By judgment I mean that which is missing when we say that someone *lacks* judgment, in the sense of not fully considering the consequences and failing to uphold the highest principles of justice and humanity and the like. Judgment is something like *phronesis*, that is, involving wisdom, prudence, even virtue.<sup>9</sup>

Why does judgment matter? The answer is again ontological. As I have been at pains to say, registrations (even nonconceptual ones) abstract from the world's detail—they approximate, do violence, privilege some things at the expense of others. That is not to deny that registrations

---

8. Two caveats. First, "judgment" is normally used in philosophy to signify an act, whereas I am using it here to name a property or capacity—something a system can *have*. Second, I am not here arguing that the notion of judgment being defended in this book is a "mark of the mental," in the sense of being a property that all human cognitive activity must manifest at all times, in order to count as human cognitive activity at all. It is important that we be able to say that someone did or said or thought something *without judgment*—by reflex, for example, or "thoughtlessly."

9. *Phronesis* traditionally concerns *practical* judgment, which is not my special concern here. Nevertheless, I take it that knowing how to act appropriately requires being grounded in concrete situations and recognizing that their subtlety and complexity may outstrip any registration of them—a theme I believe to characterize all judgment. In general, holding registration accountable requires resolute commitment to the world, not to registrations of the world—a theme that at least resonates with virtue ethics' deprecation of deontology (see, e.g., Alasdair MacIntyre, *After Virtue*, Notre Dame, IN: Notre Dame University Press, 1981).

are essential. Without them, given our finite brain (system) capacity and the world's partial disconnection, we would be at a loss, unable to orient to anything beyond our effective grasp. We would have no world.<sup>10</sup> Yet no matter their necessity, our *registrations* are not what matters—what matters is *that which we register*. To be accountable, to hold things to account, is to know the difference between the two—and to be committed to the latter, not the former.

Think of why we entrust child care to adults<sup>11</sup>—and about what can be catastrophically wrong with the response “But I did everything you said!” We want sitters to take care of our children, not of their (or our) registrations of our children. We want them to deal with any circumstances that arise, not just with those circumstances we have registered in advance. To reach out to the world in this way requires judgment, requires being an adult.<sup>12</sup> Or similarly: think of how pernicious it is to live with someone who is in love with their image (registration) of you, rather than with you “as you are.” Or again: if a driverless car makes a mistake, what is vulnerable is the person represented in their data structures—the full person of ineffably surpassing richness, worth, and detail, far transcending anything the car could register or conceive. If we are going to drive cars that can shoulder commitments, that is what we want them to be committed to: the survival of ourselves and of others, not the execution of mechanical steps that their designers once thought would be the best calculative

---

10. That is, nothing would be world *for us*—nothing would exist, for us. “The world itself,” as it were, would continue untroubled.

11. At least to emerging adults—and for that reason.

12. That is not to say that adults necessarily succeed; or that how situations are registered is irrelevant to laws and justice. That is all secondary. What *matters* is the child; it is the child who might die.

route to that end. If they are incapable of such commitments, then they are mere reckoners, and we should understand and deploy them accordingly.

If representations and registrations came with God-given guarantees of sufficiency, judgment might not be so difficult, or so important. But they do not. And they cannot. And the world is not such that they could. Ultimately, you cannot deal with the world *as* world without knowing that—without judgment.