

The Interpretation of Texts, People and Other Artifacts

DANIEL C. DENNETT
Tufts University

Judging a poem is like judging a pudding or a machine. One demands that it work. It is only because an artifact works that we infer the intention of an artificer.

W. Wimsatt and M. Beardsley¹

I want to explore four different exercises of interpretation:

- (1) the interpretation of texts (or *hermeneutics*),
- (2) the interpretation of people (otherwise known as “attribution” psychology, or cognitive or *intentional psychology*),
- (3) the interpretation of other artifacts (which I shall call *artifact hermeneutics*),
- (4) the interpretation of organism design in evolutionary biology — the controversial interpretive activity known as *adaptationism*.

My title suggests only three exercises, not four, but as we shall see when we turn to the fourth, and grounding, exercise, there is really just one exercise of interpretation here; my task is to show that these enterprises do not just have a lot in common; they are *the same project* addressed to different objects. I want to show that the canons of interpretation are the same, the problems are the same, and the illusions that beset both practitioners and their critics are the same. Finally, I will draw out some rather

¹ William Wimsatt and Monroe Beardsley, “The Intentional Fallacy,” in *The Verbal Icon*, University of Kentucky Press, 1954, p. 4.

curious implications from this examination, implications that some people find deeply satisfying and others find deeply unsettling.

1. The Interpretation of Texts

Let us begin by reminding ourselves of some of the uncontroversial facts about the interpretation of texts. Some texts are obvious, and some are relatively perplexing and inscrutable, but even the most obvious are interpretable only with the aid of rather safe, rather obvious assumptions. For instance someone who knows no English will see



TAE CAT

Figure 1.

as having the same symbol in the middle of each group, whereas we find it effortless — to the point of not noticing — to see “the cat” as the probable interpretation of this symbol string. How do we do it? We rely on the norms for the formation of letters, on the one hand, and on the other we have expectations about the likelihood that this inscription was produced with some communicative *intent*. We try to see a message in any string of letters, and in this instance we succeed. Of course *maybe* the author of this inscription *meant* to write “the cht” (or “tae cat”) — in which case we have erred in our headlong interpretation.

Our interpretation, it seems, hinges on hypotheses we have about the intent of the author of any inscription, and to a first approximation, if we are in doubt about what a bit of text means, we can seek out the author, and ask him what he meant. We do this all the time, and it works; otherwise we wouldn't do it. If the instruction book you received with your new computer says “The red light indicates that your battery is running down” and you wonder whether this means that your battery is *abnormally* running down or just running down in the course of normal operation, you can call up the author — that is, the author-of-record, the *virtual* author, the person at the other end of the 800-HELP-number — and ask him or her which interpretation was *intended*.

This wise everyday practice creates the plausible impression that whenever a question arises about the interpretation of a bit of text, one can in principle obtain the truth by asking the author, who is held to be in a privileged position to answer. Who knows better than the author what the author means? And yet the universality of this plausible principle has been famously attacked by literary critics — far from being true, it is a fallacy,

the Intentional Fallacy.² According to more or less received opinion amongst literary theorists, it is simply a mistake to suppose that the author's personal opinion about the meaning of a text is *authoritative* (as one might say!). Some theorists go so far as to assert that the author's opinion is not even a relevant datum — one might as well ask your uncle as ask the author.

These are controversial waters we are entering, and for the moment I am simply expressing — without yet endorsing — the family of claims one encounters among the hermeneuticists. Just as important as the author's reading, they say, is the reader's reading. Any reader, simply by being a reader, a *user* of the text we might say, establishes *an* interpretation of the text, and one reader's interpretation is as good as another's — including the author, who is, after all, another reader of the text he happens to have brought into the world.³ Of course some readings are better than others — that is, more interesting, more coherent, richer, less *ad hoc*, better informed — but no reading has the privilege of being singled out as *the real meaning* of the text.⁴

This extreme view is wildly implausible when the text in question is a sign saying “KEEP OFF THE GRASS” — which is virtually univocal (though hermeneuticists are perversely clever at finding alternative possible — if not plausible — readings of even the most mundane inscriptions). And when mundane equivocality arises, asking the author almost always works: we may wonder how to read Jones' shopping list, but Jones can be counted on — unless he is deranged — to know jolly well what he meant, and better than anyone else. But if the text is a poem or novel or, say, one of Plato's dialogues, the plausibility of rival interpretations is often

² Ibid.

³ Can this be as true of a spoken text as a written text? Bertrand Russell's life provides an example that suggests it can: “It was late before the two guests left and Russell was alone with Lady Ottoline. They sat talking over the fire until four in the morning. Russell, recording the event a few days later, wrote, ‘I did not know I loved you till I heard myself telling you so — for one instant I thought “Good God, what have I said?” and then I knew it was the truth’.” In *The Life of Bertrand Russell*, by R. W. Clark: London, Weidenfeld and Nicolson, 1975.

⁴ John Updike, in *Assorted Prose*, p. 204. draws our attention to a particularly inspired appropriation of one author's text by another:

On the copyright page of his first book, *The Works of Max Beerbohm*, Max found the imprint

London: JOHN LANE, *The Bodley Head* New York: CHARLES SCRIBNER'S SON
Beneath it, he wrote in pen:

This plain announcement, nicely read, Iambically runs.

. . . Indeed, were I a high priest of literature, I would have this quatrain made into an amulet, and wear it around my neck, for luck.

roughly equal, and the question of which is best — or the truth — can be wide open. In these instances, asking the author may be impossible (if the author is Plato) or in any event inconclusive. What if the author is now confused or forgetful about what he meant? Even if the text is brand new — just uttered — the author may find himself perplexed, or agnostic, about the meaning of what he uttered, or — perhaps the most titillating case — vehement and sincere in his avowal of an utterly second-rate and unconvincing interpretation. The text can indeed be better than its author realizes.

(I deliberately pass over on this occasion the debates that have raged over the relation of speaker meaning to utterance meaning, which are irrelevant to my claims here, though that perhaps is not obvious.)

There is a lurking suspicion in all such cases that however confounded the data may be, however uninvestigatable the relevant facts may be, there always *is* a “deeper fact” about what the text *really means*. Deep in the mind or brain of the author at the moment of creation lay some clue, some feature, that settles what he meant — even if he and we can gain no access to it. It is precisely this tempting suspicion that Wimsatt and Beardsley called the Intentional Fallacy. Their thesis was about the relevance of such deeper facts to the meaning of texts, not about the existence of such facts. But others have wanted to call into question the very existence of such facts. Quine’s famous thesis of the Indeterminacy of Radical Translation is the somewhat more general and radical claim that when there is inscrutability of interpretation regarding *any* human behavior — not just the behavior of text-creation — there may be no deeper facts that settle the matter.⁵

What, after all, could one find in the brain of an agent (the author of an act) that could settle between rival interpretations? The initial and entirely plausible hunch is that one might find an inner *representation* of the agent’s intention — a description of the act in question written in Mentalese. Suppose, for the moment, that there are such representations, and suppose, moreover, that they are systematic in the way a language is. Then when we uncover the representation that we expect to settle our quandary we will find that our problem recurs: now we have a new text to interpret: a string of Mentalese symbols, and who is the authority with regard to their intended meaning? Now this problem will not be serious in any case in which the inner representation “wears its meaning on its sleeve.” Just before lapsing into a coma, Paul scribbles “GRAND LEGS — SEIZE

⁵ That there are no such deeper facts is also argued for at length in my book, *The Intentional Stance*, Cambridge, Massachusetts: MIT Press/A Bradford Book, 1987, esp. pp. 37-42. See also my “Real Patterns,” forthcoming, *Journal of Philosophy*.

OURS!” on a piece of paper, and we wonder what it means. Thanks to the miracles of neurocryptology, we determine that Paul, at the time of writing, was intending to speak French. What he meant, translated now into English, was “Large legacy — sixteen bears!” and sure enough, it turns out that he had just learned that his rich uncle had bequeathed him sixteen dancing bears, not the fortune he had counted on. His suicide is explained.

But sometimes what we find inside may leave matters unimproved. Suppose, for instance, (to adapt an example from Anscombe’s pioneering essay on the topic, *Intention*⁶) that you say “Now I press Button A” while reaching out and pushing Button B. We wonder which slip you made — a slip of the tongue or a slip of the finger. Both are possible, and typically you would know which you meant to say, and do, but we can imagine a case in which you find yourself perplexed: you were under pressure, there was a case to be made for either action, you simply don’t know what you meant. Indeed just such an occasion once befell me. I was importuned to be the first base umpire in a baseball game. At the crucial moment in the game it fell to me to decide the status of the batter running to first. It was a close call, and I found myself emphatically jerking my thumb up — the signal for OUT — while yelling “SAFE!” In the ensuing tumult I was called upon to say what I had meant. I honestly couldn’t say, at least not from any privileged position. I finally decided (to myself) that since I was an unpracticed hand-signaller but competent word-speaker, my vocal act should be given the nod, but anyone else could have made just the same judgment.

In a case like this we might hope that a scan of the brain-writing would settle the issue, but suppose we find what by our neuro-cryptological lights are both intentions written in the appropriate places in the brain. When we look deeper we find an initially inscrutable set of further representations. Our epistemic position then will be really rather similar to that of the literary sleuths who have contradictory diary entries shedding rival lights on the novelist’s intended meanings. Are there deeper principles to which to appeal in settling such a case? I will return to this issue after we have had a chance to look at the other domains of interpretation.

2. Artifact Hermeneutics

Artifact hermeneutics is just textual hermeneutics applied to a wider domain. Texts are artifacts, after all. Some artifacts, such as paintings and sculptures, are notoriously — or deliciously — open to rival interpretation, and the artist, of course, is an unreliable guide. We find it much easier to believe in the inspired, possessed, entranced plastic artist who is driven

⁶ Oxford: Blackwells, second edition, 1963, p. 57.

to create works he only dimly understands (or even massively misunderstands) than we are to believe in the writer of texts who is similarly oblivious to their meaning, but the principle of authorial non-authority is the same.

Other artifacts, however, seem to have their interpretation much more securely fixed; there can be little doubt what an axe is,⁷ or what a telephone is, and what it is for; we hardly need to consult Alexander Graham Bell's biography for clues about what he had in mind. Many artifacts go out of the way to wear their meaning on their sleeves — they have labels “CAN OPENER” and “FORD ESCORT GT” printed prominently on them. They often come with user manuals or instructions — and in cases where there might be some danger of misinterpretation, their parts are prominently labeled: “THIS IS NOT A STEP” it says on the folding paint-bucket shelf of my stepladder. But in spite of these precautions, inscrutability of purpose or function is still possible in principle, and we should remember that a written label is just a text after all, as subject to quandaries of interpretation as any other design feature.




Figure 2.

Does the curiously faded label on the bottle say “FOR FEELING FAINT” or “FOR PEELING PAINT”? One had better seek further elucidation before taking a swig. More particularly, one had better examine what the substance in the bottle *can do* irrespective of what its creator intended it to do.

Archeologists and historians sometimes encounter artifacts whose meaning — whose function or purpose — is particularly obscure. It is instructive to look briefly at a few examples to see how one reasons in such cases.

⁷ For modern axes, this is certain, but prehistoric “axes” are another matter. The “Acheulean hand axe,” which was popular among our ancestors for perhaps a million years, would seem to be ill-designed for holding and chopping with. Its frisbee-like shape suggests it may have been a projectile instead. William Calvin provides an insightful discussion of this exercise of artifact hermeneutics in *The River that Flows Uphill*, San Francisco: Sierra Club Books, 1986, pp. 310, 505.



(Copyright © Professor Derek de S. Price, Yale University)

Figure 3

The Antikythera mechanism, discovered in 1900 in a shipwreck, and dating from ancient Greece, is an astonishingly complex assembly of bronze gears. What was it for? Was it a clock? Was it the machinery for moving an automaton statue, like Vaucanson's marvels of the 17th century? It

was — almost certainly — an orrery or planetarium, and the proof of that is that it would be a *good* orrery. That is, calculations of the periods of rotation of its wheels led to an interpretation that would have made it an accurate (Ptolemaic) representation of what was then known about the motions of the planets.


The great architectural historian Viollet-le-Duc described an object called a *cerce*, used somehow in the construction of cathedral vaults. He hypothesized that it was a movable piece of staging, but a later interpreter, John Fitchen, argued that this could not have been its function (see figure 4, taken from Fitchen's classic, *The Construction of Gothic Cathedrals: A Study of Medieval Vault Erection* [Chicago: The University of Chicago Press, 1981]; © 1961 by John Fitchen). For one thing, the *cerce* would not have been strong enough in its extended position, and, as the figure shows, its use would have created irregularities in the vault webbing which are not to be found. Fitchen's extended and elaborate argument concludes that the *cerce* was no more than an adjustable template, a conclusion he supports by coming up with a much more elegant and versatile solution to the problem of temporary support of web courses.

The important feature in these arguments is the reliance on optimality considerations; it counts against the hypothesis that something is a cherry-pitter, for instance, if it would have been a demonstrably inferior cherry-pitter.⁸

Occasionally, an artifact loses its original function and takes on a new one. People buy old-fashioned sad-irons not to iron their clothes with, but to use as bookends or doorstops; a handsome jam pot can become a pencil holder, and lobster traps get recycled as outdoor planters. The fact is that sad-irons are much better as bookends than they are at ironing clothes — at least compared to the competition today. And a Dec-10 mainframe computer today makes a nifty heavy-duty anchor for a large boat mooring. No artifact is immune from such appropriation, and however clearly its *original* purpose may be read from its current form, its new purpose may be related to that original purpose by mere historic accident — the fellow who owned the obsolete mainframe needed an anchor badly, and opportunistically pressed it into service.

Inventors of artifacts are no more immune to confusion than authors of texts. It is possible that someone setting out with every intention of creating a new kind of alarm clock succeeds, in spite of himself, in creating

⁸ The example comes from my earlier discussion of artifact hermeneutics, in "Beyond Belief," in Andrew Woodfield, ed., *Thought and Object*, Oxford: Clarendon Press, 1982, pp. 41-43. (pp. 155-56 as reprinted in *The Intentional Stance*. See further pp. 304, 311-12.)



*Viollet-le-Duc's Cerce Device as Support for Each Web Course
During the Erection of the Vault*

The smaller-scale drawing shows a cerce, based on Viollet-le-Duc's representation and description. Its extended position clearly indicates how one slotted board laps the other. Hung vertically as support for the stones of a web course, it is seen (in the detailed section) that the stones of any given course cannot line up throughout: those that lean against the far board (shown in outline) tilt much more than those that lean against the near board (shown hatched). As no such break does occur in the alignment of the web stone-coursing, it is obvious that the cerce device was not used in this fashion, in spite of Viollet-le-Duc's assertion that it was.

Figure 4.

something that can *charitably* be described as merely a new kind of paper-weight. Consider how the Intentional Fallacy looks when applied to artifacts: the inventor is not the final arbiter of what an artifact is, or is for; the *users* decide that. The inventor is just another user, only circumstantially and defeasibly privileged in his knowledge of the functions and uses of his device. If others can find better uses for it, his intentions, clearheaded or muddled, are of *mere historical interest*. That is, it may indeed be an incontrovertible historical fact that a certain artifact was created by someone with a particular purpose very clearly represented — both in his head, in Mentalese, and in written “specs” and blueprints, we may suppose — but this historical fact, while it establishes something about how the artifact was intended at the outset, *may* shed no valuable light on the functions it can and does actually serve. Turing and Von Neumann might disavow the use of their computing machine as a word-processor, but so what? It is a perfectly marvelous word-processing machine, and many of the features of its most recent mutations are clearly designed with that end in mind.

Of particular interest for our purposes is the interpretation of robots. These machines, often designed to preserve themselves in the face of environmental challenges, do indeed perform “actions” controlled by internal representations specifying their purposes. Suppose we face the sort of issue that confronted our human button-pusher — the robot performs an ambiguous or equivocal act and we wonder what it meant to do. There are several avenues we can explore: (1) we can ask what, by our lights (as robot users) it *ought to have meant*, and such use of a principle of charity or optimality may quickly settle the issue. Or (2), if no stable interpretation is forthcoming by that route — something seems to have gone wrong — we can turn to the designers of the robot to find out what they had in mind. They may not tell us, or may have lost the blueprints and other specifications, and what they tell us may be out of date in another way: the design of the robot may have shifted in response to its environment in ways they didn’t foresee and can no longer decipher. So their opinions, while valuable evidence, are not guaranteed to settle any outstanding question of interpretation we may have. The robot may have taken on a somewhat autonomous interpretation. When we reconsider the robot in the context of the costs and benefits of its current environment, we may arrive at a better interpretation of its internal states than its original designers can muster.

3. Interpreting People and other Organisms

Now it is time to turn to the postponed question: when there is residual indeterminacy or controversy about the interpretation of a bit of human

psychology, how can we settle it? When we *adopt the intentional stance* toward a person, we use an assumption of rationality or cognitive/conative optimality to structure our interpretation, but when something goes wrong — when we find evidence of apparent sub-optimality or breakdown, the problems of interpretation multiply swiftly. We can no longer reside much faith in the agent's own opinions, in many instances, and so, if the actions in question mean anything at all, it is not anything for which we have much direct evidence. What should we do? Consult the designers — just as we do with other artifacts. For we are artifacts, after all, designed by natural selection to provide reliable survival vehicles for our genes into the indefinite future. (The remainder of this section draws on material that appears in “Interpreting Monkeys, Theorists, and Genes,” in *The Intentional Stance*, pp. 269-86.)

How can we consult Mother Nature? By adopting the intentional stance toward the process of natural selection itself and looking for the somewhat covert design rationales of the features we discover. In short, we can become adaptationists.⁹

In biology, the adaptationists assume optimality of design in the organisms they study, and this practice is viewed askance by some other biologists, since it seems to them to invoke a doctrinaire optimism. Why should anyone today suppose that an organism, just because it has evolved, is in any way optimally designed? There is now a mountain of evidence and good theory in population genetics, for instance, to show that under many conditions maladaptations are fixed, developmental constraints limit phenotypic plasticity, etc. But this challenge is ill-posed; the critics who remind the adaptationists of these complications are already talking past the opposition.

We take on optimality assumptions not because we naively think that evolution has made this the best of all possible worlds, but because we must be interpreters if we are to make any progress at all, and interpretation requires the invocation of optimality. The adaptationist strategy in biology seeks to answer “why?”-questions in exactly the same way the intentional strategy in psychology does. Why, asks the folk psychologist, did John decline the invitation to the party? The presumption is that there is a (good) reason — at least in John's eyes. Why, asks the adaptationist, do these birds lay four eggs? The adaptationist starts with the supposition that there is a (good) reason: that four eggs is better, somehow, than two

⁹ For more on adaptationism and the strategy of adopting the intentional stance towards evolution, see *The Intentional Stance*, chapters 7 and 8. The present paper is designed to provide the context without which the brief remarks on interpretation in those chapters tend to be misapprehended — or dismissed, via a generous but unwanted invocation of the principle of charity!

over the term still to be found in many biology textbooks: “pre-adaptation,” with its entirely spurious suggestions of foresight or preparation.

These cases have been the focus of considerable controversy in biology, and some biologists, peering into the abyss that has opened, are tempted to renounce *all* talk of function and purpose. They are right about one thing: there is no stable intermediate position.¹³ If you are prepared to make any claims about the function of biological entities — for instance, if you want to maintain that it is perfectly respectable to say that eyes are for seeing, and the eagle’s wings for flying — then you take on a commitment to the principle that natural *selection* is well-named. In Elliot Sober’s terms, there is not just selection *of* features but selection *for* features.¹⁴ And without this “discriminating” prowess of natural selection, we would not be able to sustain functional interpretations at all.

Certainly we can describe all processes of natural selection without appeal to such intentional language, but at enormous cost of cumbersome, lack of generality, and unwanted detail. We would miss the pattern that was there, the pattern that permits prediction and supports counterfactuals. The “why” questions we can ask about the engineering of our robot, which have answers that allude to the conscious, deliberate, explicit reasonings of the engineers (in most cases) have their parallels when the topic is organisms and their “engineering.” If we work out the rationales of these bits of organic genius, we will be left having to attribute — but not in any mysterious way — an emergent appreciation or recognition of those rationales to natural selection itself.

How can natural selection do this without intelligence? It does not consciously seek out these rationales, but when it stumbles on them, the brute requirements of replication ensure that it “recognizes” their value. The illusion of intelligence is created because of our limited perspective on the process; evolution may well have tried all the “stupid moves” in addition to the “smart moves,” but the stupid moves, being failures, disappeared from view. All we see is the unbroken string of triumphs.¹⁵ When we set

¹³ “Among evolutionary biologists, there are those who condemn the identification of anatomical structures as having specific adaptational significance, on the ground that such structures do not face selection individually, but only in the company of the rest of the organism. This makes ascriptions of adaptational ‘content’ to a part of the organism indeterminate, since a different ascription together with other adjustments in our adaptational identifications can result in the same level of fitness for the whole organism. In the philosophy of psychology, the dual of this thesis is reflected in the indeterminacy of interpretation.” Alexander Rosenberg, “Intentional Psychology and Evolutionary Biology (Part I: The Uneasy Analogy),” *Behaviorism*, 14, pp. 15-27.

¹⁴ Elliot Sober, *The Nature of Selection*, 1984, Cambridge, Massachusetts: MIT Press/A Bradford Book.

¹⁵ See also Jonathan Schull, “Are Species Intelligent?” forthcoming in *Behavioral and Brain Sciences*, for a fascinating argument to show that the process of natural selection

NovaGene, a biotechnology company in Houston, has adopted the policy of “DNA branding”: writing the nearest codon rendering of their company trademark in amino acids in the “extra” or “junk” DNA of their products (according to the standard abbreviations, asparagine-glutamine-valine-alanine-glycine-glutamic acid-asparagine-glutamic acid © NQVAGENE).¹⁷ This suggests a new exercise in radical translation for philosophers: how could we confirm or disconfirm the hypothesis that trademarks — or service manuals — were discernible in the bulk of DNA that is apparently uninvolved in the direction of the formation of the phenotype? Richard Dawkins’ gene’s-eye view predicts, and hence could explain, the presence of this meaningless “selfish DNA,”¹⁸ but that doesn’t show that it *couldn’t* have a more dramatic source — and hence a meaning after all.

Would a closer look at the designs reveal some telltale discontinuities? Natural selection, lacking foresight, cannot see the wisdom in *reculer pour mieux sauter* — step back in order to jump forward better. If there are designs that cannot be approached by a gradual, step-wise re-design process in which each step is at least no worse for the gene’s survival chances than its predecessor, then the existence of such a design in nature would seem to require, at some point in its ancestry, a helping hand from a foresightful designer — either a gene-splicer, or a breeder who somehow preserved the necessary succession of intermediate backsliders until they could yield their sought progeny. But couldn’t such a leap forward — a “saltation” in the terms of evolutionary theorists — be a mere lucky leap? At what point do we dismiss the hypothesis of cosmic accident as too improbable, and settle for the hypothesis of interventionist engineers?¹⁹

These questions suggest — but do not prove, of course — that there may be no foolproof marks of natural (as opposed to artificial) selection. Indeed, all the biologists I have queried on this point have agreed with me that there are no sure marks of artificial selection. Would this conclusion, if proven, be a terrible embarrassment to the evolutionists in their struggle against creationists? One can imagine the hubbub: “Scientists Concede: Darwinian Theory Cannot Disprove Intelligent Design!” But this is to mistake the status of orthodox evolutionary theory. It would be most foolhardy for any defender of the theory of natural selection to claim that it gives one the power to read history so finely from present data as to rule

¹⁷ *Scientific American*, June 1986, p. 70-71.

¹⁸ Richard Dawkins, *The Extended Phenotype*, 1982, chapter 9: “Selfish DNA, Jumping Genes, and a Lamarckian Scare.”

¹⁹ See the discussions of gradualism, saltation and probability in R. Dawkins, *The Blind Watchmaker*, London: Longman Scientific, 1986.

out, absolutely, the earlier historical presence of rational designers. It may be a wildly implausible fantasy, but it is a possibility after all.

In our world today, there are organisms we *know* to be the result of foresighted, goal-seeking redesign efforts, but that knowledge depends on our direct knowledge of recent historical events (we've actually watched the breeders at work), but these special events might not cast any fossil shadows into the future. To take a simpler variation on our thought experiment, suppose we were to send Martian biologists a laying hen, a Pekingese dog, a barn swallow and a cheetah and ask them to determine which designs bore the mark of intervention by artificial selectors. What could they rely on? How would they argue? They might note that the hen did not care "properly" for her eggs; some varieties of hen have had their instinct for broodiness bred right out of them and would soon become extinct were it not for the environment of artificial incubators human beings have provided for them. They might note that the Pekingese was pathetically ill-suited for fending for itself in any demanding environment. The barn swallow's fondness for carpentered nest sites might fool them into the view that it was some sort of pet, and whatever features of the cheetah convinced them that it was a creature of the wild might also be found in greyhounds and have been patiently encouraged by breeders. Artificial environments are themselves a part of nature, after all.

Prehistoric fiddling by intergalactic visitors with the DNA of earthly species cannot be ruled out, except on grounds that it is an entirely gratuitous fantasy. Nothing we have found (so far) on earth so much as hints that such a hypothesis is worth further exploration. (And note — I hasten to add, lest creationists take heart — that even if we were to discover and translate such a "trademark message" in our spare DNA, this would do nothing to rescind the claim of the theory of natural selection to explain all design in nature without invocation of a foresighted Designer-Creator *outside the system*. If the theory of evolution by natural selection can account for the existence of the people at NovaGene who dreamt up DNA branding, it can also account for the existence of any predecessors who may have left their signatures around for us to discover.) The power of the theory of natural selection is not the power to prove exactly how (pre-)history was, but only the power to prove how it could have been, given what we know about how things are.

Adaptationist thinking, then, may often be unable to answer particular questions about specific features of the historical mechanisms, the actual etiology, of a natural design development, even while it can succeed in formulating, and even confirming — insofar as confirmation is ever possible — a functional analysis of the design. The difference between a design's

having a free-floating (unrepresented) rationale in its ancestry and its having a represented rationale may well be indiscernible in the features of the design, but this uncertainty is independent of the confirmation of that rationale for that design. Moreover, the historical facts about the process of design development, even when we can discover them, are equally neutral when we move in the other direction: they are unable to resolve questions about the rationale of the design on which our interpretation of its activities depends. We should still hope science will eventually uncover the historical truth about these etiological details, but not because it will resolve all our Aristotelian “why” questions, even when they are cautiously and appropriately posed.

It is time to take stock of this survey of varieties of interpretation. To begin, we can fill in a chart with the examples discussed.

Varieties of Interpretation

	“Obvious”	Equivocal	“Exapted”
TEXTS	KEEP OFF THE GRASS telephone book	“feeling faint”	Beerbohm poem
ARTIFACTS	axe telephone	Antikythera mechanism cerce	iron bookends word processor
EVOLVED	eye wing	Stegosaurus fins turkey wattles	panda’s thumb
PEOPLE’S INTENTIONAL STATES AND ACTS	“pass the sugar” believing it’s raining	“I now push Button A” “Out!” [Safe]	“I love you.” —Bertrand Russell

It is not just that each exercise of interpretation yields cases with the same features. In every case, or at least so I have argued, the principles of interpretation are the same. While the preponderant sorts of evidence in the different varieties are quite different, there are no sources of evidence or styles of argument that are available in principle for one but not another variety of interpretation. In each instance, the “user” is no more privileged than the author. And even though the user “selects” a design of current “utility,” utility can change swiftly, and the user is not omniscient. We may well not be able to project the user’s rationale into the future —

any more than the original author's. So what something is *really for now* is no more authoritatively fixed by the current user's "intentions" than by any other intentions.

Curiously, then, we can get *better* grounds for making reliable functional attributions (functional attributions that are likely to continue to be valuable interpretation-aids in the future) when we ignore

- (1) "what the author says" in the case of texts
- (2) the "R and D" record in the case of other artifacts
- (3) the fossil record in the case of evolved features
- (4) what people avow in the case of intentional states.

There is no "privileged access" to be had in any quarter. The intention, if any, with which an item was originally introduced determines, at most, what function the author hoped or intended the item to serve, and that hope too is a product of interpretation. We cannot begin to make sense of functional attributions until we abandon the idea that there has to be one, determinate, *right* answer to the question: What is it for? And if there is no deeper fact that could settle that question, there can be no deeper fact to settle its twin: What does it mean?