

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Representation: Where Philosophy Goes When It Dies.

Permalink

<https://escholarship.org/uc/item/36b2p9zz>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Author

Slezak, Peter

Publication Date

2001

Peer reviewed

Representation: Where Philosophy Goes When It Dies.

Peter Slezak (p.slezak@unsw.edu.au)

Program in Cognitive Science, University of New South Wales
Sydney NSW 2052 AUSTRALIA

Abstract

Robert Cummins (1996, p.1) has characterised the problem of mental representation as “*the* topic in the philosophy of mind for some time now”. This remark is something of an understatement. The same topic was central to the famous controversy between Nicolas Malebranche and Antoine Arnauld in the Seventeenth Century and remained central to the entire philosophical tradition of “ideas” in the writings of Locke, Berkeley, Hume and Kant. I show that the recurrence of certain deep perplexities about the mind is a systematic and pervasive pattern, confirming Jerry Fodor’s disparaging remark: “Cognitive science is where philosophy goes when it dies” (Fodor, 1994b, p. 110).

The Tripartite Schema

Recently Bechtel (1998, p. 299) states the essentials of a modern theory of representation: “There are ... three interrelated components in a representational story: what is represented, the representation, and the user of the representation”.

Z: System Using Y → Y: Representation → X: Thing Represented

Among the problematic assumptions, Bechtel’s diagram and discussion crucially fail to distinguish internal and external representations. Bechtel’s conception in this regard is not idiosyncratic but almost universal in cognitive science (Newell 1986, p. 33). As we will note presently, in the case of pictorial images, the assimilation of internal and external representations tacitly encourages the illegitimate postulate of a user or external observer - the notorious homunculus. The same tacit assimilation of external and internal representations is at the heart of Searle’s “refutation” of symbolic AI and also leads to the doctrine that we think “in” language (Carruthers 1996, Slezak forthcoming a). The assimilation just noted in Bechtel will also be seen in the seemingly unrelated problem of consciousness and the mind-body problem (Place 1956), *inter alia*.

The tripartite scheme appears obvious and innocuous enough but Bechtel’s diagram (modified here) is a variant of the scheme which we see throughout the long history of the subject. Thus, for example, nothing could seem more remote from modern theories in cognitive science today than Malebranche’s (1712/1997) seventeenth century doctrine of “the vision

of all things in God”. On the contrary, however, despite the theological trappings, it is instructive to recognize the profound affinity of Malebranche’s views with those at the very forefront of theorising today in psychology and artificial intelligence: Malebranche’s theory is just Bechtel’s tripartite model (Nadler 1992), and the modern problem of representation is how to avoid the notorious difficulties clearly articulated by his critic Arnauld (1683/1990).

It is no accident that Gibson’s ‘ecological’ approach involves a direct realism which has been proposed as alternative to the representationalism of computational theories. This is merely one form in which the Malebranche-Arnauld debate is being rehearsed today. This celebrated debate is described by Nadler (1989) as a debate between an ‘object theory’ of ideas and an ‘act theory’, respectively. He explains

... the object theory of ideas involves a commitment to a representationalist or *indirect* realist theory of perception, such as Malebranche (and, on the traditional reading, Locke) put forth. An act theory of ideas, on the other hand, forms the core of Arnauld’s perceptual direct realism. If ideas are representational mental acts [rather than entities], then they can put the mind in direct cognitive contact with the world - no intervening proxy, no *tertium quid*, gets in the way. (1989, p.6)

It is no accident that recent proponents of ‘situated cognition’ have been complaining of exactly the same indirect, mediated conception in computational theories of cognition. For example, J. Greeno (1989) unwittingly echoes Arnauld:

I am persuaded ... we are connected directly with the environment, rather than connected indirectly through cognitive representations.

... An individual in ordinary circumstances is considered as interacting with the structures of situations directly, rather than constructing representations and interacting with the representations. (1989, p. 290)

Precursors: Pointless Exercise?

Despite the skepticism expressed by Gaukroger (1996), precursors of modern cognitive science provide an independent, extensive source of insight into contemporary issues and, conversely, are themselves elucidated in novel ways unavailable to traditional

scholarship (Yolton 1996, Slezak 1999, 2000). The possibility of two-way elucidation arises from the extraordinary persistence of the seemingly simple problem of saying “in some illuminating way, what it is for something in the mind to represent something” (Cummins 1996, p.1). From Yolton’s (1984, 1996) statement of earlier concerns we see their parallel with the contemporary problem: Scholastics’ notions of ‘intelligible species’ and Cartesian talk of ideas were striving for some way to explain the conformity or agreement between ideas and objects.

If Malebranche and Arnauld anticipated contemporary concerns about representation in cognitive science, then it is clear that the current theoretical problem has nothing to do with the theoretical framework of symbolic, computational approaches as universally assumed. Indeed, the recurrence of essentially the same dispute in widely varying contexts today suggests that the underlying problem does not arise essentially from the special features of any one of them. I suggest that we may discern the same underlying problem at the heart of notorious disputes such as ‘The Imagery Debate’, Searle’s Chinese Room conundrum, the thinking-in-language debate, ‘situated cognition’ and a number of others which have been prominent and recalcitrant.

No Representations?

The ‘Cognitive Revolution’ was characterized by a re-discovery of the indispensability of internal representations following their repudiation by Skinnerian behaviourism. There is considerable irony in recent approaches which appear to reject internal representations once again (Brooks 1991, Freeman and Skarda 1990, Clark and Toribio 1994, Greeno 1989, van Gelder, 1998). Notwithstanding Eliasmith’s (1996) claim, these views are not plausibly seen as a return to behaviourism, but they are symptoms of the profound difficulties posed by the phenomena. It is sobering to notice that Arnauld’s critique of Malebranche exactly prefigures these recent attacks on representational theories. Arnauld’s treatise *On True and False Ideas* is concerned to repudiate what he describes as “imaginary representations”, saying “I can, I believe, show the falsity of the hypothesis of *representations*” (1683/1990, p.77) for “one must not make use of alleged *entities* of which we have no clear and distinct idea in order to explain the effects of nature, whether corporeal or spiritual” (1683/1990, p. 65).

Tables & Chairs: Bumping Into Things

Fodor (1985a) joked that philosophers are notorious for having been prey to absurd, eccentric worries such as the “fear that there is something fundamentally unsound about tables and chairs”. Nevertheless, he optimistically opined that sometimes “mere” philosophical worries turn out to be *real* as in the case of the representational character of cognition. However, far from being a

contrast with the traditional anxiety about tables and chairs, modern scientific disputes concerning representations appear to be *identical* with this notorious worry!

Thus, it is surely no accident that, reflecting upon Fodor’s (1980) ‘methodological solipsism’, Jackendoff (1992, 161) asks facetiously “Why, if our understanding has no direct access to the real world, aren’t we always bumping into things?” Though intending a mild parody, Jackendoff captures precisely the paradox charged against Locke and also Malebranche, who Nadler (1992, p. 7) says “is often portrayed ... as enclosing the mind in a “palace of ideas,” forever cut off from any kind of cognitive or perceptual contact with the material world”. Thus, Jackendoff’s satire is evocative of Samuel Johnson’s famous refutation of Berkeley’s “ingenious sophistry” by kicking a stone. Of course, Berkeley’s “sophistry” is just the worry about the reality of tables and chairs.

Imagery: The Pictorial Theory

The ‘Imagery Debate’ is perhaps the most remarkable modern duplication of seventeenth century controversies. In this re-enactment, among the *dramatis personae* Pylyshyn plays Arnauld against Kosslyn’s Malebranche. Significantly, the central error identified by Arnauld of ascribing corporeal properties to mental ones exactly the one charged by Pylyshyn (1973, 1981) against Kosslyn and related to Cummins’ (1996) point that internal representations do not function by being understood”.

Kosslyn’s (1994) pictorial account of imagery takes mental images to represent by virtue of a relation of resemblance to their objects and by virtue of actually *having* spatial properties which they represent. Furthermore, “depictive” representations in a “visual buffer” are taken to have the specific function of permitting a re-inspection of images by the higher visual apparatus. Not surprisingly, this “quasi-perceptual” model has been repeatedly charged with the error of importing an ‘homunculus’. The charge is vigorously rejected on the grounds that “the theory is realized in a computer program” (Kosslyn, Pinker, Smith & Schwartz, 1979, p. 574), but undischarged homunculi can lurk in computational models just as easily as in traditional discursive theories (see Slezak 1992, 1994, 1995, 1999). Thus, Kosslyn, Sokolov and Chen (1989) offer a diagram of the visual imagery system which is a profusion of inter-connected boxes and arrows. The box labeled “visual buffer” contains another box labeled “attention window” which is left unexplained. This box is, in fact, the observer in the ‘theater’ which is the source of the traditional problem. The elaborate diagram is reducible to the same tripartite schema we have seen in Malebranche. Significantly, following Descartes, Arnauld explicitly pointed to the seductive error of taking pictures as an appropriate model of mental representation (Arnauld 1683/1990, p.

67) and he cites the *camera obscura* as an erroneous model for imagery. Thus, retinotopic maps on the visual cortex cited by Kosslyn (1994), p. 14) as vindicating the pictorial theory are pictures alright, but only for the *theorist*.

Descartes Déjà Vu.

A related unlikely indication of the relevance of early philosophy to current problems is seen in Edelman's (1998) work on perception. Despite its concern with the latest theories of perception, the central problem is stated in terms identical with that of the entire tradition of writers on 'ideas'. Edelman writes: "Advanced perceptual systems are faced with the problem of securing a principled (ideally, veridical) relationship between the world and its internal representation." Edelman's solution "is a call for the representation of similarity instead of representation by similarity". This might have been taken verbatim from Descartes's *Treatise of Man* or *Dioptrics* where he said "the problem is to know simply how [images] can enable the soul to have sensory perceptions of all the various qualities of the objects to which they correspond - not to know how they can resemble these objects" (Descartes 1985, 1, 165).

Mind-Body Problem

The pervasive error seen starkly in Kosslyn's TV screen metaphor reveals the link between the various problems in cognitive science and the traditional mind-body problem. In the classic statement of materialism, U.T. Place (1956) argued that the rejection of materialism is based on the qualitative features of subjective experience. Although these features have recently been supposed to constitute the "hard" problem of consciousness (Chalmers 1996), Place suggested that they are the source of the 'phenomenological fallacy'. This is "the mistake of supposing that when the subject describes his experience, how things look, sound, smell, taste, or feel to him, he is describing the literal properties of objects and events on a particular sort of internal cinema or television screen." Place's diagnosis has been revived and given prominence by Dennett (1991), however the fallacy would be more aptly named the 'Malebranchian Theater'.

Thinking In Language

Just as we seem to be looking at pictures when we imagine visually, so we appear to talk to ourselves when we think. Indeed, Carruthers (1996) who seeks to revive what he acknowledges to be an unfashionable doctrine explicitly bases his argument on such evidence of introspection. This is the evidence that we sometimes find ourselves in a silent monologue, talking to ourselves *sotto voce*, as it were.

However, in a neglected article, Ryle (1968) suggested that the very idea that we might think "in"

language is unintelligible, and the undeniable experience of talking to ourselves cannot support any claim about the vehicles of thought. It is significant that Ryle mentions *en passant* among the equally problematical cases, that in which we claim to see things in our 'mind's eye' - taken to involve mental pictures of some kind. Ryle's comparison and his warning is unwittingly confirmed by Carruthers (1996, 1998) who explicitly invokes Kosslyn's pictorial account of imagery as support for his own analogous theory. In doing so, however, Carruthers only brings into relief the notorious difficulties of his own model which relies on representations - sentences of natural language - which are, like pictures, paradigmatically the kind requiring an external intelligent observer.

Connectionism & Cognitive Architecture

The crucial difference between symbolic and connectionist architectures is said to be the absence of *explicit* representations with constituent structure (Fodor & Pylyshyn 1988, Fodor & McLaughlin 1990). This issue may turn precisely on the same question we have seen. Specifically, the distinction between explicit and implicit representation appears to be based on a tacit appeal to the criterion of intelligibility or discernability to an external observer. D. Kirsh (1990, p. 340) points out that in connectionist systems, "it is becoming increasingly difficult ... to track the trajectory of informational states these mechanisms generate. There is no doubt that we must find some method of tracking them; otherwise there is no reason to think of them as more than complex causal systems." Explicitness is characterised as a matter of "directly reading off" information from "visible structures" and the "immediate grasp" of information which is "directly available" or "immediately readable" (1990, p. 356), but the obvious question is: By whom? In an influential article Ramsey, Stich and Garon (1991) question traditional folk psychology because "in many connectionist networks it is not possible to localize propositional representation beyond the input layer" (1991, p. 209). Again we may ask, By whom? Significantly, Ramsey et al. imply that the difficulty of identifying representations in a neural net is "a real inconvenience to the connectionist model builder" (1991, p. 209). However, as Cummins' (1996, p.102) warns, "Internal representations are not exploited by being understood" by the programmer".

Symbols & Searle

Searle's (1980) Chinese Room conundrum appears to have an identical logical structure to those I have noted. In this case, a crucial equivocation on distinct meanings of 'meaning' has led to the postulation of symbols having meaning in an *observer-relative* sense in which a representation is necessarily apprehended and understood by someone. However, intelligibility *to the theorist* must be irrelevant to Searle's question of whether a

system has genuine, 'original' intentionality (see Slezak 1994, 1999). As before, rejecting this inappropriate criterion of meaning actually amounts to rejecting a certain conception of representations or, equivalently, rejecting the *agent* as homunculus in the system. Intentionality and the directness of cognition is achieved, following Arnould, by eliminating a conception of symbols as intermediate objects to be apprehended as if they were external representations.

Searle's (1980) criterion for judging intentionality in his Chinese Room amounts to Carruthers' (1996) claim that the language of thought or 'mentalese' must be English, since the symbols are to be understood by a fully comprehending intelligent person. Accordingly, the much-discussed conundrum is best understood, not as a challenge to 'strong AI' as such, but as a *reductio ad absurdum* of symbols conceived as being intelligible to an observer. It is significant that this mistake is not Searle's alone, for it is implicit in the orthodox computational view of cognitive science which sees its origin in Frege-Russell formal symbols requiring an interpretation (Newell and Simon 1976, Newell, 1986). In AI, too, this conception has been explicitly embraced by Nilsson (1987, 1991) and embodied in the CYC program of Lenat and Feigenbaum which has been caricatured by Smith (1991) as the 'Electric Encyclopedia'. In these cases the question of meaning of mental representations is confused between whether representations are *intelligible* and whether they are *explainable*.

Searle's conundrum is evoked by Glanvill's response in 1661 to Descartes' coding theory of perception. "But how is it, and by what Art doth the soul read that such an image or stroke in matter ... signifies such an object? Did we learn such an Alphabet in our Embryo-state?" (quoted in Yolton 1984, p. 28). Echoing Searle, Glanvill suggests that the "motions of the filaments of nerves" learn the quality of objects by analogy with the way in which a person learns to understand a language, for otherwise "the soul would be like an infant who hears sounds or sees lips move but has no understanding of what the sounds or movements signify, or like an illiterate person who sees letters but 'knows not what they mean'" (1984, p. 28).

Logicism & Observer Attribution

Within AI, an independent, though parallel, debate has been proceeding about the classical symbolic conception, the "logicist" view, according to which an abstract formal system gets its meaning from a model theory - that is, the intended interpretations of the designer (Newell and Simon 1976, Nilsson 1987, 1991). Woods (1987), Smith (1987), Rosenschein (1985) and others have argued that this classical logical view is fundamentally misguided in its conception of the way in which a system gets to relate to the external world by embracing the specific conception of "observer attribution" (see Hadley 1995). Birnbaum (1991, p. 62)

says that AI mistakenly adopts a theory from logic which is quite inappropriate to capture what it means for a system to have beliefs.

Misrepresentation

The problem of misrepresentation has arisen for causal or co-variation theories of intentional content (Dretske 1996, Fodor 1994a) since these theories seem to be unable to capture the way a mismatch might arise between a representation and the world. If a mentalese token 'mouse' might be caused not only by mice but also by shrews, then the symbol must *ipso facto* mean 'shrew' and cannot be in error.

The puzzle might be accounted for by noting that it arises from tacitly adopting the stance of external interpreter: The very problem itself cannot be coherently formulated except in terms of judgements which are not part of the scientific, explanatory enterprise. The veridicality of representations is not a property which can play any role in the functioning of representations or the explanation of them. Like the picture on a jigsaw puzzle, the meaning of representations conceived as semantically evaluable in this way is for our own benefit and not intrinsic to the arrangements of interlocking components.

The very concern with misrepresentation arises from tacitly adopting a questionable assumption endorsed by Davidson (1975) that having a belief requires also having the *concept* of belief, including the concept of error. However, it seems that animals might have beliefs even if they are unable to know that they have them and reflect on their truth value. A cat can surely be correct in thinking that a mouse is in a certain hole without having the *concepts* of belief and truth.

Argument From Illusion

The modern problem of misrepresentation is a variant of the classical 'argument from illusion' employed in support of Locke's 'ideas' and A.J. Ayer's (1940) sense-data as the immediate objects of perception. The parallel should not be surprising since an illusion in the relevant sense is precisely a misrepresentation. The traditional argument, just like Dretske's and Fodor's, turns on the possibility of a mismatch between mental representations and their referents in the external world.

Responding to Ayer, Austin (1962, p.61) remarked on the "curious" and "melancholy fact" that Ayer's position echoes that of Berkeley. Of course, this is the same melancholy fact that Fodor's "real" problems of representation are identical with the traditional concerns about the reality of tables and chairs. Questions of veridicality for ideas and sense-data arose from precisely the same assumptions as Fodor's - namely, the spurious possibility of a comparison between representations and the world. Twin Earth puzzles, too, seem to be an unnoticed variant on the problem of misrepresentation (Slezak, forthcoming b).

It should be less surprising that the classical arguments for 'ideas' should be akin to the modern case for representations when it is noticed that the 'argument from illusion' is effectively an 'argument from imagery'. The proverbial illusory pink elephant as the immediate object of perception is a visual image *par excellence*.

Illusion of the Intelligent Reader

We might expect a compelling kind of error to emerge in unrelated domains of theorising about the mind. Chomsky has drawn attention to the way in which traditional grammars produce an illusion of explanatory completeness while, in fact, they have "serious limitations so far as linguistic science is concerned" (Chomsky 1962, p 528). The success of the grammar depends on being "paired with an intelligent and comprehending reader". Here we see an entirely different version of the homunculus problem. Chomsky notes that in judging the adequacy of traditional grammars the unnoticed reliance on the user's linguistic ability is illegitimate because it is just what the theory is supposed to explain (Chomsky, 1962, p.528). Evidently, this is just the issue captured in Cummins' distinction, in a different context, between 'meaning' and 'meaningfor' (1996, p. 86) and is evidently the problem also for pictorial images or thinking in natural language.

Conclusion

Fodor (1968, p. vii) once remarked: "I think many philosophers secretly harbor the view that there is something deeply (ie. conceptually) wrong with psychology, but that a philosopher with a little training in the techniques of linguistic analysis and a free afternoon could straighten it out." Thirty years later, the suspicion of deep conceptual problems at the heart of philosophy and psychology is more clearly justified. By adopting a broader perspective we may see why the sorry fortunes of the two disciplines have been inextricably linked.

References

- Arnould, A. (1683/1990). *On True and False Ideas*. Trans. S. Gaukroger, Manchester University Press.
- Austin, J.L. (1962). *Sense and Sensibilia*. Oxford University Press.
- Ayer, A.J. (1940). *The Foundations of Empirical Knowledge*. Macmillan.
- Bechtel, W. (1998). Representations and Cognitive Explanations. *Cognitive Science*, 22, 3, 295-318.
- Birnbaum, L. (1991). Rigor Mortis: A Response to Nilsson's 'Logic and Artificial Intelligence'. *Artificial Intelligence*, 47, 57-77.
- Brooks, R. (1991). Intelligence Without Representation. *Artificial Intelligence*, 47, 139-159.
- Carruthers, P. (1996). *Language, Thought and Consciousness*. Cambridge University Press.
- Carruthers, P. (1998). Thinking in Language? P. Carruthers and J. Boucher eds., *Language and Thought*. Cambridge University Press.
- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.
- Chomsky, N. (1962). Explanatory Models in Linguistics. In E. Nagel. P. Suppes & A. Tarski, eds., *Logic, Methodology and Philosophy of Science*. Stanford University Press, 528-550.
- Clark, A. & Toribio, J. (1994). Doing Without Representing? *Synthese*, 101, 3, 401-431.
- Cummins, R. (1996). *Representations, Targets and Attitudes*. Bradford/MIT Press.
- Davidson, D. (1975). Thought and Talk. In S. Guttenplan ed., *Mind and Language*. Clarendon Press.
- Dennett, D.C. (1991). *Consciousness Explained*. Penguin.
- Descartes, R. (1985). *The Philosophical Writings of Descartes*, in 2 Volumes. Translated by J. Cottingham, R. Stoothoff & D. Murdoch, Cambridge University Press.
- Dretske, F. (1986). Misrepresentation. In S. Stich & T. Warfield eds., *Mental Representation*. Blackwell 1994.
- Edelman, S. (1998). Representation is Representation of Similarities. *Behavioral and Brain Sciences*, 21, 449-498.
- Eliasmith, C. (1996). The third contender. *Philosophical Psychology*, 9, 4, 441-463.
- Fodor, J.A. (1968). *Psychological Explanation*. Random House.
- Fodor, J.A. (1980). Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavioral and Brain Science*, 3, 63-109.
- Fodor, J.A. (1985a). Presentation. In B.H. Partee, S. Peters and R. Thomason eds., *Report of Workshop on Information and Representation*. NSF System Development Foundation, 106-117.
- Fodor, J.A. (1985b). Fodor's Guide to Mental Representation. In *A Theory of Content and Other Essays*. MIT Press, 1990, 3-29.
- Fodor, J.A. (1994a). *The Elm and the Expert*, MIT Press.
- Fodor, J.A. (1994b). Concepts: A Potboiler. *Cognition*, 50, 95-113.
- Fodor, J.A. & McLaughlin, B. (1990). Connectionism and the Problem of Systematicity. *Cognition*, 35, 183-204.
- Fodor, J.A. & Pylyshyn, Z. (1988). Connectionism and Cognitive Architecture. *Cognition*, 28, 3-71.
- Freeman, W.J. & Skarda, C.A. (1990). Representations: Who Needs Them? In J. L. McGaugh, N. Weinberger & G. Lynch eds. *Brain Organization and Memory Cells*. Oxford Univ Press.
- Gaukroger, S. (1996). *Descartes: An Intellectual Biography*. Oxford University Press.

- Greeno, J.G. (1989). Situations, Mental Models and Generative Knowledge. In D. Klahr and K. Kotovsky eds. *Complex Information Processing: The Impact of Herbert A. Simon*. Lawrence Erlbaum.
- Hadley, R.F. (1995). The Explicit- Implicit Distinction. *Minds and Machines*, 5, 219-242.
- Jackendoff, R. (1992). *Languages of the Mind*. Bradford/MIT Press.
- Kirsh, D. (1990). When is Information Explicitly Represented? In P. Hanson, ed. *Information, Language and Cognition*. Univ of British Columbia Press.
- Kosslyn, S.M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. MIT Press.
- Kosslyn, S.M., M.A. Sokolov & J.C. Chen 1989. The Lateralization of BRIAN. In D. Klahr & K. Kotovsky eds. *Complex Information Processing: The Impact of Herbert A. Simon*. Lawrence Erlbaum.
- Kosslyn, S., Pinker, S., Smith, G. & Schwartz, S. (1979). On the demystification of mental imagery. *The Behavioral and Brain Sciences*, 2, 535-581.
- Malebranche, N. (1712/1997). *The Search After Truth*. Trans. T.M. Lennon & P.J. Olscamp. Cambridge University Press.
- Nadler, S. (1989). *Arnauld and the Cartesian Philosophy of Ideas*. Manchester University Press.
- Nadler, S. (1992). *Malebranche and Ideas*. Oxford University Press.
- Newell, A. (1986). The Symbol Level and the Knowledge Level. In Z. Pylyshyn and W. Demopoulos eds. *Meaning and Cognitive Structure*. Ablex.
- Newell, A. & Simon, H.A. (1976). Computer Science as Empirical Inquiry. *Communications of the ACM*, 19, 113-126.
- Nilsson, N.J. (1987). Commentary on McDermott. *Computational Intelligence*, 3, 202-203.
- Nilsson, N.J. (1991). Logic and Artificial Intelligence. *Artificial Intelligence*, 47, 31-56.
- Place, U.T. (1956). Is Consciousness A Brain Process?. In J. O'Connor ed., *Modern Materialism: Readings on Mind-Body Identity*. Harcourt Brace.
- Pylyshyn, Z. (1973). What the Mind's Eye Tells the Mind's Brain. *Psychological Bulletin*, 80, 1, 1-24.
- Pylyshyn, Z. (1981). The Imagery Debate. In N. Block, ed. *Imagery*. MIT Press.
- Ramsey, W., Stich, S. and Garon, J. (1991). Connectionism, Eliminativism and the Future of Folk Psychology. In W. Ramsey, S. Stich and D. Rumelhart, eds. *Philosophy and Connectionist Theory*. Lawrence Erlbaum.
- Rosenschein, S.J. (1985). Formal Theories of Knowledge in AI and Robotics. *New Generation Computing*, 3, 345-357.
- Ryle, G. (1968). A Puzzling Element in the Notion of Thinking. In P.F. Strawson ed., *Studies in the Philosophy of Thought and Action*. Oxford University Press, 7-23.
- Searle, J. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3, 417-424.
- Slezak, P. (1992). When Can Images Be Reinterpreted. *Proceedings of 14th Conference of the Society for Cognitive Science*. Lawrence Erlbaum, 124-129.
- Slezak, P. (1994). Situated Cognition: Empirical Issue, Paradigm Shift or Conceptual Confusion. *Proceedings of 16th Conference of the Society for Cognitive Science*. Lawrence Erlbaum.
- Slezak, P. (1995). The Philosophical Case Against Visual Imagery. In P. Slezak, T. Caelli and R. Clark eds. *Perspectives on Cognitive Science*. Ablex.
- Slezak, P. (1999). Situated Cognition: Empirical Issue, Paradigm Shift or Conceptual Confusion? In J. Wiles & T. Dartnall eds. *Perspectives on Cognitive Science, Vol. 2*, Ablex.
- Slezak P. (2000). Descartes' Startling Doctrine of the Reverse-Sign Relation. In S. Gaukroger, J. Schuster, J. Sutton eds. *Descartes' Natural Philosophy*. Routledge, 542-556.
- Slezak, P. (forthcoming a). Thinking About Thinking, *Language and Communication*.
- Slezak, P. (forthcoming b). Representing. In P. Staines, H. Clapin & P. Slezak eds. *Representation in Mind*. Greenwood.
- Smith, B.C. (1987). The Correspondence Continuum. *CSLI Report 87-71*.
- Smith, B.C. (1991). The Owl and the Electric Encyclopedia. *Artificial Intelligence*, 47, 251-288.
- van Gelder, T. (1998). The Dynamical Hypothesis in Cognitive Science. *Behavioral and Brain Sciences*, 21, 615-665.
- Woods, W.A. (1987). Don't Blame the Tools. *Computational Intelligence*, 3, 228-237.
- Yolton, J.W. (1984). *Perceptual Acquaintance from Descartes to Reid*. University of Minnesota Press.
- Yolton, J.W. (1996). *Perception and Reality*: Cornell University Press.