Where Syllogistic Reasoning Happens
An Argument for the Extended Mind Hypothesis

Georg Theiner (gtheiner@indiana.edu)
Department of Philosophy, 1033 E. Third St.
Bloomington, IN 47405 USA

Abstract

Does cognition sometimes literally extend into the extra-organismic environment (Clark, 2003), or is it always "merely" environmentally embedded (Rupert, 2004)? Underlying this current border dispute is the question about how to individuate cognitive processes on principled grounds. Based on recent evidence about the active role of representation selection and construction in learning how to reason (Stenning, 2002), I raise the question: what makes two distinct, modality-specific pen-and-paper manipulations of external representations – diagrams versus sentences – cognitive processes of the same kind, e.g. episodes of *syllogistic reasoning*? In response, I defend a "division of labor" hypothesis, according to which external representations are dependent on perceptually grounded neural representations and mechanisms to guide our behavior; these internal mechanisms, however, are dependent on external representations to have their syllogistic content fixed. Only their joint contributions qualify the extended computational process as an episode of *syllogistic reasoning* in good standing.

Keywords: Philosophy of Psychology; Extended Mind Hypothesis; Situated Cognition; Diagrammatic Reasoning

Extended or Merely *Embedded* Cognition?

In recent years, the development and use of epistemic artifacts – tools for thinking – have come to play a central role in the evolutionary and psychological explanation of human intellectual prowess. We deploy them to access new kinds of information (e.g. compass, sextant); to off-load, re-represent, and return information (e.g. pen and paper, writing systems); and to introduce new active sources of information-processing (e.g. abacus, pocket calculator, neural implants). They allow us to accomplish cognitive tasks that would otherwise remain beyond the ken of our "naked" biological brains.

How should we characterize these intriguing capacity-enhancing interactions? Proponents of the *Extended Mind Hypothesis* (EMH) view them as episodes of cognitive processing which literally stretch beyond the head to include elements of the organism's local environment among its proper parts (Clark & Chalmers, 1998; Clark, 2003; Wilson & Clark *forthcoming*). From the perspective of a more conservative competitor, which Rupert (2004) calls the *Hypothesis of Embedded Cognition* (HEMC), cognitive processes are safely encased within the brain, but sometimes "depend very heavily, in hitherto unexpected ways, on organismically external props and devices" (p. 393).

Underlying this border dispute is a question about how to individuate cognitive systems. According to EMH, what constitutes a single, extended cognitive system is determined by the level of duration, reliability, and functional integration of the coupling between biological on-board and environmental off-board resources assembled for the performance of specific cognitive tasks (Wilson & Clark, *forthcoming*). In return, critics such as Adams and Aizawa (*forthcoming*) have branded the "coupling-constitution fallacy" as "the most common mistake extended mind theorists make" (p.2) when they infer, on the basis of showing that an external resource is coupled in some principled fashion to a cognitive agent, that this resource constitutes a mereological *part* of the agent's cognitive apparatus. Contested by their criticism is a tacit premise which would close the suggested logical gap:

The A-B PRINCIPLE: In certain cases where an understanding of A's relation to B is significantly relevant to our understanding of A, we should posit a single system, A-B, as a single unit of study[1].

As a theoretical hypothesis in cognitive science, EMH should be evaluated on a case-by-case basis in terms of its empirical and methodological credentials. Consequently, if we can enumerate conditions under which psychologists are profitably committed to specific applications of the A-B Principle for demarcating the boundaries of a cognitive system, we have good reasons to prefer EMH over HEMC as a theoretical gloss of what's going on in these cases. In this paper, my defense of the A-B Principle within the psychology of deductive reasoning is predicated on showing two things: first, that it is befitting to characterize pen-and-paper performances of syllogistic problem-solving as products of a *single* hybrid computational system realized by parts of the brain, body, and elements of the locally scaffolded environment; second, which insights we can only obtain by treating the extended system as a single *cognitive* system.

The Argument from Modality

The dominant *naturalist* approaches to the psychology of so-called "mental" (i.e. in-the-head) syllogistic reasoning exhibit various reflexes of the traditional belief that logical competence is a fixed part of human nature. For instance, it is widely assumed that human subjects are essentially homogeneous in what kind of internal representation system they use; that the main theoretical role of external

[1] A stronger, unrestricted version of this principle is discussed and rightly criticized by Rupert (2004).

formalisms is to analyze these already pre-established systems; and that learning how to reason is a matter of tapping into a universal, biologically predisposed inferential mechanism. These fundamental assumptions are shared by otherwise competing theories of the cognitive architecture supposed to underlie our basic inferential capacities, such as Mental Models (Johnson-Laird 1983), Mental Logics (Rips 1994), Euler Circles (Stenning & Oberlander 1995), or the innate modules of evolutionary psychology (Cosmides 1989). Surprisingly little attention has been paid to how subjects actually go about finding and using external representations to solve syllogistic problems, and individual differences in their ability to employ sentential versus diagrammatic modalities.

In *Seeing Reason*, Stenning (2002) has outlined an alternative *constructivist* approach that highlights the active roles of representation selection and meta-representational knowledge in learning how to reason, as well as the cognitive re-organization that occurs within a learner upon exposure to a specific formalism. His findings suggest that "[…] there is more than one mental representational approach to syllogisms" (p.128), and that the received assumption of homogeneity is no more warranted in the psychology of internal representations than it is in logic, which recognizes a wide variety of external representation systems. Instead, Stenning advocates the need for a comparative theory of the cognitive effects of representational modalities.

As Stenning points out (p.17f), it is clear that such a theory logically requires that there is some relation of invariance under which the compared cognitive processes are *semantically* equivalent, in the sense that they all involve the representation of the same logical properties and relations characterizing the syllogistic target domain. Within naturalist frameworks, the presumed semantic invariance of variegated acts of cognition can be explained in terms of a universal class of "distilled" syllogistic representations in the brain on which the dedicated inferential processor gets to operate after all extraneous, modality-specific features of external information displays have been stripped away. But if there is no single, universal representational mechanism in the brain which underlies our manifold attempts to solve syllogisms, as Stenning's research indicates, *what, if anything, about our cognitive machinery explains why they are all instances of syllogistic reasoning*? A standardized presentation of this meta-theoretical "problem of modality" will facilitate our subsequent discussion.

Suppose that HOLIST and SERIALIST[2] belong to different student populations whose modality preferences systematically vary with respect to their background skills and cognitive styles. HOLIST gets taught to use Euler Circles (EC) which suit her reasoning style, while SERIALIST gets taught to use Natural Deduction[3] (ND) which suits her distinct style. When they are given syllogisms (in English) to solve using pen and paper, HOLIST deploys a modality-specific computational mechanism M1 which operates on EC, utilizing a class IR1 of internal representations; SERIALIST deploys a different modality-specific computational mechanism M2 which operates on ND, utilizing a distinct class IR2 of internal representations. M1 is an agglomerative mechanism which is differentially sensitive to topological and quasi-mechanical features of EC, but incapable of processing information presented in ND. Likewise, M2 is a discursive mechanism which is differentially sensitive to the syntactically mediated concatenative features of ND, but incapable of processing information presented in EC. What makes our supposition that SERIALIST and HOLIST both engage in *syllogistic reasoning* true?

Providing an affirmative solution to our problem of modality suggests the following (compressed) argument for EMH:

1. SERIALIST and HOLIST both engage in syllogistic reasoning.
2. That SERIALIST and HOLIST both engage in syllogistic reasoning is best explained by the fact that they manipulate mental representations which have the same syllogistic content.
3. The only relevantly available representations are {IR1, EC} for HOLIST, and {IR2, ND} for SERIALIST.
4. Tokens of IR1 and IR2 are semantically type-distinct when considered in isolation.

Therefore, the mental representations which have the same syllogistic content are external tokens of EC and ND[4].

I shall take our supposition in premise (1) at face value, and suggest granting premise (3) at least for the sake of the argument. The remainder of this paper is dedicated to a defense of the critical premises (2) and (4).

On the Domain-Identity of Cognitive Processes

Premise (2) calls upon the content of mental representations to explain the domain-identity of cognitive processes. It is based on the principle that the domain D of a cognitive process carried out by some psychological mechanism M is determined by the informational content of the mental representations on which M operates. A process of syllogistic reasoning must thus involve the execution of a mechanism operating on mental representations capable of tracking the relation of logical consequence for a well-defined sub-domain of monadic first-order logic.

The best defense of this principle is to consider an alternative behavioral criterion based purely on input-output equivalence. Suppose we held that two cognitive processes

[2] See Pask (1975). Monaghan and Stenning (1998) report strikingly symmetrical aptitude-by-treatment interactions between HOLISTS and SERIALISTS in their assessment of teaching/learning-effectiveness with diagrammatic versus sentential methods; see also Monaghan et al. (1999).

[3] Stenning and Yule (1997) present a case-based algorithm for syllogistic reasoning which uses only the sentential fragment of the natural deduction calculus.

[4] For the time they are actually utilized to solve a syllogistic task.

are domain-identical just in case they both map the same classes of sensory inputs on the same classes of behavioral outputs. Such a criterion requires that we be able to define the relevant classes of inputs and outputs on logically independent grounds without referring to mental representations. One way this could in principle be done is by using the Ramsey-Carnap-Lewis method (Lewis, 1972) for defining theoretical terms. It employs a meta-language in which classes of external stimuli and behavioral responses are treated as observationally primitive, and mental states are interdefined functionally as states that causally interact with each other and sensory stimuli to produce observable behavior. To see whether this behavioral criterion defeats the present argument for EMH, let us reflect on what the relevant "inputs" and "outputs" would be.

According to a first proposal, we take the class of syllogisms in English as our inputs, and either 'NVC'[5], or 'VC' plus conclusions in English, as our outputs. If M1 and M2 are input-output equivalent in this sense, the behavioral criterion does indeed entail that HOLIST and SERIALIST are both reasoning syllogistically, but also validates EMH by definition. For both subjects, the hypothesized sub-processes of forward-translating the input into external formalisms, their manipulation, and the backwards-translation play an integral causal role in computing the overall input-to-output mappings. But the behavioral criterion decrees under which conditions two causal processes are domain-identical *qua* being cognitive. Hence any internal *or external* intermediary states that are causally relevant parts of the mechanism deployed to achieve the desired mappings would appear as Ramseyfied theoretical terms in our psychological theory, and therefore be part of the cognitive process in question.

It is presumably in recognition of this very consequence that critics of EMH, in particular Adams and Aizawa (2001), have rejected the largely behavioral criterion of the cognitive at play in the *Parity Principle* invoked by Clark and Chalmers (1998). Meant to offset our vernacular prejudice in favor of the skull as a theoretically relevant boundary to delimitate cognitive processes, Parity says that if, as we confront some task, a part of the world functions as a process which, *were* it to go on in the brain, we would accept as part of the cognitive process, then that part of the world *is* (for that time) part of the cognitive process. Far from implying any deep functional similarity between inner and outer processes, Parity is equally compatible with their often disparate, but complementary and synergetic contributions in producing behavioral outputs. It only requires that the overall behavioral profile of an extended cognitive system displays enough of the central features and dynamics of brain-bound cognitive systems (Clark 2003; 2005). A purely behavioral criterion for the domain-identity of cognitive processes should thus be fairly unattractive for the HEMC-theorist on independent grounds.

[5] 'VC': 'valid conclusion', 'NVC': 'no valid conclusion'.

As a second proposal, the HEMC-theorist might suggest to break down the overall input-output mapping into a series of simpler mappings[6], but insist that during each step, certain inputs are "caught" by the brain, cognized only in the brain, and then "tossed" back to paper as outputs. If taken literally, this catch-and-toss-redescription clearly does not help to solve our current predicament. By hypothesis, M2 is not differentially sensitive to graphical features of EC, hence M2 could not compute *any* of the required mappings after translation; *mutatis mutandis* for M1. It follows that HOLIST and SERIALIST would not engage in cognitive processes of the same domain.

This problem can be remedied by defining semantically relevant equivalence classes between specific operations afforded by EC and ND. For instance, the case-based method of ND presented by Stenning and Yule (1997) is apt to reveal representational homologies between sentential and diagrammatic methods in their ability to track *critical individuals*, i.e. individuals which are fully determinate with regard to all three properties. This is an important feature to share, because any syllogistically valid inference can be captured by a single critical individual.

However, the resulting criterion of domain-identity would thus be representational and not behavioral, as we can see from applying Marr's (1982) three-level computationalist framework for understanding cognition. The main goal of a level-1 analysis of the behavioral task to be performed is to pin down a precise input-output function. Giving a computational specification of an external representation system rather corresponds to a level-2 analysis, because it provides explicit schemes for representing syllogistic inputs and outputs, together with a sequence of mechanical steps to carry out the relevant subtasks. But in Marr's framework, level-2 explanations of information-processing systems traffic in *mental* representations. Hence if we rely on level-2 analyses of the beyond-the-head parts of M1 and M2 to ground their semantic domain-identity, the second proposal also ends up supporting EMH. I conclude that the HEMC-theorist is well-advised to accept some stronger representational criterion of cognitive domain-identity, similar to the one proposed in premise (2).

Division of Labor among Inner and Outer

When mental representations are posited as theoretical entities in psychological explanations of behavior, they are expected to do double-duty: first, to be semantically evaluable, i.e. to serve as stand-ins which are *about* or *refer* to certain things; second, to play a causal role in the determination of behavior. These two requirements constrain the kinds of states apt to play the role of mental representations. Considering *aboutness* as a true mark of the

[6] For the case of HOLIST, this decomposition might look like this: *input_1* = class of syllogisms in English, *output_1* = two separate premise diagrams in EC; *input _2* = two separate premise diagrams in EC, *output_2* = agglomerated premise diagram in EC; *input _3* = combined premise diagram in EC, *output_3* = 'NVC', or 'VC' plus conclusion in English.

mental, some critics of EMH have argued that external representations are never cognitive because their semantic content is always merely derived from the representational properties of neurally realized mental states (Adams & Aizawa, 2001). Challenging this intuition, I maintain that neither internal nor external resources by themselves have the right functional profile to satisfy both requirements at once. External representations are dependent on perceptual-motor and neural mechanisms to guide our behavior. But as I suggest in premise (4), these internal mechanisms are dependent on external representations to have their syllogistic content fixed. Only their joint contributions qualify the extended causal process as an episode of *syllogistic reasoning*; hence the A-B Principle applies.

I shall offer two sources of evidence in support of my division-of-labor hypothesis: for the computational part, I highlight some features of EC as a graphical algorithm; for the cognitive part, I then take a look at some of the leading theories of mental content-determination.

Syllogistic in-the-head reasoning is hard because of the high task-demands it places on working memory. It requires the temporary binding of properties into specifications of individuals, grouping these individuals into models, and often comparing several alternative models. These bindings have to be effected anew for each task, without reliance on previously established bindings stored in long-term memory. The direct semantic interpretability of spatial relations in EC dovetails nicely with the problem-solving capacities of HOLIST's resource-limited brain to facilitate the processibility of syllogistic inferences (Stenning & Oberlander, 1995).

First, taking spatial containment of a point in a curve to denote set-membership, the topology of EC enforce a fully determinate representation of an individual's properties, since every point in a plane is either inside or outside any closed curve. Assigning mnemonic linguistic tags to indicate that distinct circles denote distinct properties, EC provide perceptually stable representations of the required bindings which are easy to attend to. Assuming that all closed curves are continuous and refer uniquely to properties, EC are *self-consistent* insofar no single diagram can represent an inconsistent set of propositions.

Second, to distinguish individuals that *must* exist from those that only *may* exist, HOLIST uses the cross-notation as an "abstraction trick" to represent only minimal models of each premise. To combine two premise diagrams, HOLIST registers the B-circles, and chooses the arrangement between A- and C-circles which creates the most subregions. If a cross-marked region is bisected during this agglomeration, it is excised. A persisting cross the final diagram indicates a valid inference. Since we further know that syllogistic logic is *case-identifiable* in the sense that there is only a single critical individual (i.e. a unique minimal model) on which any valid syllogistic inference depends, EC are well-suited to make this meta-logical property of syllogisms particularly transparent to a HOLIST learner.

Third, by allowing the brain to treat circles as "hoops" and crosses as "nails", EC acts as a mechanical device that mirrors logical constraints imposed by the premises. For instance, if a nail prevents the A- and C-circles from being pulled apart, a positive conclusion follows; if a nail prevents a complete alignment, a negative conclusion can be inferred. HOLIST can thus bring her prior mechanical knowledge to bear on syllogistic tasks. Fourth, comparing the meta-logical properties of EC and ND, we can show that they are both sound and complete *as syllogistic proof procedures*, and that there are close representational homologies between the modality-specific operations afforded by the two formalisms (Stenning & Yule, 1997).

In sum, EC are well-designed epistemic tools because their geometrical and mechanical properties effectively mirror the logical constraints of the syllogistic domain. They act as external computational resources which allow HOLIST's brain to accomplish syllogistic feats by performing much simpler visuo-spatial tasks. But the semantic content of the cognitive processes underpinning these visuo-spatial tasks is not inherently syllogistic in nature, as I shall now argue. Considered in the context of our argument, this means that the only relevantly available mental representations with intrinsic syllogistic content must be EC, even though they are located outside the organism. As a case in point, let us consider the cross-marked region in a final EC-diagram as a putative mental representation ER of the unique minimal model (UMM) capturing a syllogistically valid inference.

According to Dretske's (1988) indicator account of mental content, ER represents UMM because it has an acquired causal role within the behavioral economy of the larger computational system (e.g. the presence of ER leads HOLIST to emit 'VC', its absence leads to an 'NVC'-response), and it acquired that function because it reliably indicated the presence of UMM. According to a modified[7] teleo-semantic account (Millikan, 1984; Papineau, 1987), ER represents UMM since EC have been culturally selected by HOLISTS as their preferred syllogistic representation-producing device *because* ER indicates UMM. According to Cummins' (1996) structuralist account, ER represents UMM because it structurally mirrors the three-property specification of UMM, and this isomorphism is exploited for behavioral control.

EC bring about their desired behavioral effects only relative to a fixed, ecologically normal backdrop which includes an organism's on-board capacities to construct and manipulate external symbols. But this kind of causal context-dependence is in fact common to many familiar examples of "encodings". It equally applies to neural structures prescribing bodily motions, genes affecting developmental outcomes, or even pieces of C++ code controlling traffic lights. Whenever we single out specific parts of an extended process as the relevant causal difference-makers coding for a certain outcome of that

[7] The modification concerns the requirement that EC have been selected as a result of cultural instead of biological evolution.

process, we implicitly assume (and/or ensure) that the requisite reading environment for the code is in fact in place, such that *ceteris* are indeed *paribus* (Clark, 1998).

This does not make inner representations simple duplicates of outer representations, though. Stenning (2002) notes that "at the very least one would expect external representations to obviate the need for some internal memory representations, replacing them by perceptual-motor processes with their own ephemeral representations" (p.124). Following Barsalou (1999), I take the relevant inner vehicles to be perceptually grounded working memory representations that reflect the brain's previous sensory and proprioceptive experience of interacting with external objects. Generated to specify efficient behavioral routines for an embodied agent's pen-and-paper manipulations, the neural *second-order* representations are perceptually fine-tuned to various modality-specific aspects of external *first-order* representations which do not have any syllogistic significance at all. According to my conjecture, these inner second-order representations are "action-oriented" (Clark, 1997) representations which depict outer first-order representations partly in terms of typical manipulations to be performed on these outer representations.

For instance, the precise shape, distance and positional placement of an A-circle vis-à-vis some B-circle *on paper* is logically insignificant as long as (say) A is properly contained in B. However, the *neural* representation of their exact spatial configuration might determine whether it is easier for a subject to "move" the first premise diagram on the second diagram when she redraws their compounded diagram, or *vice versa*. In the course of reasoning, M1 deploys neural tokens representing "PIN-PREVENTS-SMALLER-RIGHT-HOOP-FROM-BEING-HORIZONTALLY-PULLED-APART-FROM-UNDERLYING-BIGGER-LEFT-HOOP", while M2 deploys neural tokens representing "SHAPE-OF-SMALLISH-*B*-SPACED-PRETTY-CLOSE-TO-'→'-MATCHES-SHAPE-OF-MUCH-BIGGER-*B*-IN-LINE-BELOW". The domain-generality of these perceptually grounded inner vehicles is beneficial from a computational point of view, because subjects can partly co-opt pre-existing skills when they learn how to manipulate novel formalisms. But from a semantic perspective, it reveals our current theoretical predicament. If we abide by our representational criterion of individuating cognitive processes, the within-the-head parts of M1 and M2 are two distinct perceptual-motor processes, and not two instances of syllogistic reasoning.

We can bring out the informational asymmetry between ER and IR *qua* syllogistic representations more clearly if we employ the notion of a *proprietary domain*. The proprietary domain D of a cognitive mechanism M is the type and range of information M has been designed[8] to process, because it involves the manipulation of mental representations whose proper function is to carry information about properties of D (Millikan, 1984). For our purposes, I shall leave it open whether the design of the cognitive mechanism is a result of biological evolution, learning, or cultural evolution. According to this definition, the within-the-head part of M1 is not proprietary to the syllogistic domain. We acquire our cognitive abilities to understand overlaps, detect bisections, and recognize constraints on the movement of physical objects in the context of carrying out a motley of quite unrelated spatial and mechanical tasks. Accordingly, the syllogistic content of IR during pen-and-paper manipulations is only *derived* from the syllogistic content of ER. Consider, once again, the same theories of mental content to which we have appealed before. In Dretske's terms, IR acquired the causal role within M1 to detect UMM *only* because ER reliably indicates the presence of UMM. In teleo-semantic terms, IR has acquired the function of representing UMM within M1 *only* since EC have been culturally selected by HOLISTS as their preferred syllogistic representation-producing device *because* ER indicates UMM.

But the proprietary domain of a cognitive mechanism is defined in terms of the *intrinsic* semantic content of the mental representations on which it operates. Therefore, it is only because we consider the within-the-head part of M1 together with its beyond-the-head part to constitute a single, extended cognitive system, that the instantiation of M1 becomes a cognitive process of *syllogistic reasoning* in good standing; the same sort of considerations apply to M2. That HOLIST and SERIALIST both engage in syllogistic reasoning is therefore best explained by the fact that they manipulate semantically type-identical mental representations (EC and ND) located outside their respective heads.

Conclusions

In this paper, I have presented an argument meant to serve as a "tie-breaker" in favor of EMH over HEMC. It is based on a meta-theoretical problem in the psychology of reasoning: what makes two distinct, modality-specific pen-and-paper-manipulations of external representations – diagrams versus sentences – cognitive processes of the same kind, namely episodes of *syllogistic reasoning*? As a solution, I have proposed a "division of labor" hypothesis, according to which external representations are dependent on perceptually grounded neural representations and mechanisms to guide our behavior; these internal mechanisms, however, are dependent on external representations to have their syllogistic content fixed. Only their joint contributions qualify the extended computational processes as episodes of *syllogistic reasoning* in good standing. I argue that my solution constitutes an explanatory context where psychology is profitably committed to the A-B Principle, because it requires that we consider organismic and environmental resources as part of a single, extended

[8] I shall leave it open whether the design of the cognitive mechanism is a result of biological evolution, learning, or cultural evolution.

cognitive system. It follows that cognition sometimes literally extends beyond the head into the world.

My analysis also sheds some new light on a residual naturalist bias underlying the original formulation of Parity, in which a process completely carried out in the brain is assumed to provide our default intuition of a cognitive process, which is then used to set a bar for any external process that also ought to count as such. But according to our constructivist approach, many of our most prominent cognitive technologies start their life outside the human brain, and only gradually disappear behind the skull as a result of intensive learning and repeated practice. Doesn't my argument entail that an expert's fully internalized manipulations of EC aren't episodes of syllogistic reasoning in good standing, because the deployed neural representations have their syllogistic content in a derived manner only? To block that inference, we should also apply Parity backwards: if, as we confront some task, a part of the *brain* functions as a process which, were it done using *environmental* resources, we would accept as part of the cognitive process, then that part of the brain is part of the cognitive process.

Acknowledgments

This research has been supported by the Robert H. Ferrell Endowed Fellowship and a Louise McNutt Graduate Fellowship which have been awarded to me by the College of Arts and Sciences at Indiana University Bloomington. I have benefited greatly from discussions with Colin Allen, Jonathan Weinberg, Robert Goldstone, Kenneth Aizawa, and David Landy about extending the mind.

References

Adams, F., & Aizawa, K. (2001). The Bounds of Cognition. *Philosophical Psychology, 14*, 43-64.

Adams, F., & Aizawa, K. (*forthcoming*). Defending the Bounds of Cognition. In R. Menary (Ed.), *The Extended Mind*. Aldershot, Hants: Ashgate.

Barsalou, L. W. (1999). Perceptual Symbol Systems. *Behavioral and Brain Sciences, 22*, 577-609.

Clark, A. (1997). *Being there: putting brain, body, and world together again*. Cambridge, Mass.: MIT Press.

Clark, A. (1998). Twisted Tales: Causal Complexity and Cognitive Scientific Explanation. *Minds and Machines, 8*, 79-99.

Clark, A. (2003). *Natural-born cyborgs: minds, technologies, and the future of human intelligence*. New York: Oxford University Press.

Clark, A. (2005). Intrinsic content, active memory and the extended mind. *Analysis, 65*, 1-11.

Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis, 58*(1), 7-19.

Cosmides, L. (1989). The logic of social exchange: has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition, 31*, 187-276.

Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, Mass.: MIT Press.

Dretske, F. I. (1988). *Explaining behavior: reasons in a world of causes*. Cambridge, Mass.: MIT Press.

Johnson-Laird, P. N. (1983). *Mental models: towards a cognitive science of language, inference, and consciousness*. Cambridge, Mass.: Harvard University Press.

Lewis, D. (1972). Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy, 50*, 249-258.

Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.

Millikan, R. G. (1984). *Language, thought, and other biological categories: new foundations for realism*. Cambridge, Mass.: MIT Press.

Monaghan, P., & Stenning, K. (1998). *Effects of representational modality and thinking style on learning to solve reasoning problems.* Proceedings of at the 20th Annual Conference of the Cognitive Science Society, 716-21. Lawrence Erlbaum Associates, Maweh, NJ.

Monaghan, P., Stenning, K., Oberlander, J., & Sönstrod, C. (1999). *Integrating Psychometric and Computational Approaches to Individual Differences in Multimodal Reasoning*. Paper presented at the 21st Annual Cognitive Science Society Conference, 405-10. Lawrence Erlbaum Associates, Maweh, NJ.

Papineau, D. (1987). *Reality and representation*. New York: B. Blackwell.

Pask, G. (1975). *Conversation, cognition and learning: a cybernetic theory and methodology*. Amsterdam; New York: Elsevier.

Rips, L. J. (1994). *The psychology of proof: deductive reasoning in human thinking*. Cambridge, Mass.: MIT Press.

Rupert, R. (2004). Challenges to the Hypothesis of Extended Cognition. *The Journal of Philosophy, 101*, 389-428.

Stenning, K. (2002). *Seeing reason: image and language in learning to think*. Oxford ; New York: Oxford University Press.

Stenning, K., & Oberlander, J. (1995). A Cognitive Theory of Graphical and Linguistic Reasoning: Logic and Implementation. *Cognitive Science, 19*, 97-140.

Stenning, K., & Yule, P. (1997). Image and language in human reasoning: a syllogistic illustration. *Cognitive Psychology, 34*, 109-159.

Wilson, R. A., & Clark, A. (*forthcoming*). How to Situate Cognition: Letting Nature Take its Course. In M. Aydede & P. Robbins (Eds.), *The Cambridge Handbook of Situated Cognition*.