

Visual Attention, Conceptual Content and Doing it Right

In this paper, I argue that by understanding how vision is directly involved in action, we can see that in the case I shall consider, action on *artifacts*, vision has conceptual content. Thus, the whole-sale denial of conceptuality to vision is false. I arrive at this conclusion by arguing that there is a link between memory and concepts that provides an empirical angle on the question of conceptual content. I then present a novel account of the link between intention and perception and discuss what role concepts might play in vision and how their involvement therein amounts to a form of visual attention. Finally, I give some empirical evidence supporting the claim that visual attention has conceptual content.

This essay concerns three questions about visual perception:

- (1) Is visual content conceptual?
- (2) What does it mean to say that a concept is involved in vision?
- (3) How is vision involved in action?

I claim that by first answering (3), at least in the case of action on *artifacts*, we can provide at least partial but informative answers to the other two. Understanding how action works advances our understanding of aspects of vision.

Consider (1), perhaps the most controversial. Among the arguments in this area are those that focus on *perceptual function* which assume that the role of perception in various domains sets constraints on the nature of perceptual content. Specifically, philosophers have focused on the role of perception as providing materials for cognition (viz. belief). I shall call this view of perception's role vis-à-vis belief the *input-view*.¹

The relevance of perceptual function to the debate concerning conceptual content has been largely guided by the input-view, and thus, practically speaking, many philosophers have adhered to an *input-only* view of perceptual function, passing over other roles perception plays.² Indeed, this emphasis on the perception-belief connection has guided our understanding of the relation between perception and action. The orthodox causal approaches to action, descending from Davidson's work (1963), treat action as the ultimate product of a causal chain originating with perception leading to a belief/desire pair to intention and then to action. This causal process is constituted by events (and perhaps occurrent states), and the causal influence of perception on action is always indirect or *mediated*. There is always an intermediate event, the proximate cause of action, typically intention, which intervenes between perception and action. It is worth

¹ See Hurley 1998.

² Some exceptions: Cussins, 1990 (also Evans 1982) takes perceptual content relevant to action to be non-conceptual. Cf. Campbell 2002; Noë, 2004.

asking, however, what reasons we have for accepting this account of the perception-action link. In fact, a correct account sees perception as more directly tied to action in a way I shall bring out.³ Along the way, we shall uncover a novel argument for the claim that vision has conceptual content.

Some philosophers, however, find (1) confusing if not confused (e.g., Stalnaker, 2003). This suggests that the notion of *conceptual content* requires elucidation. Accordingly, I begin with two clarifications: first, I distinguish two notions of conceptual content: *state-conceptualism* and *content-conceptualism*.⁴ I then bring issues in the psychology of concepts to bear on the philosophical debate, deploying the notion of *semantic memory* as a bridge. I argue that semantic memory functions as a necessary basis for conceptual thought. Consequently, we can investigate the involvement of concepts via the involvement of semantic memory, and this provides an empirical angle on (1). I shall then argue that vision must be directly involved in action and explain how concepts function in the visual guidance of action on artifacts. In answering (2), I suggest that concepts facilitate visual *focusing* on and *parsing* of the artifact to be acted on and that these are phenomena of visual attention. I then argue provisionally, based on empirical work on semantic memory and action, that concepts are involved in our visually attending to the world. Admittedly, this conclusion is tied to the vicissitudes of empirical findings and subsequent research may demand reconsideration. The key point, however, is to establish a link between philosophical and psychological concerns as a way to advance the philosophical question about conceptual content.

³ On argument for this is presented below in questioning whether intentions could be the proximate causes of action. I present other reasons for doubting the standard picture (Wu, ms.).

⁴ The distinction is discussed by Stalnaker, op. cit., Heck, 2000, Byrne 2004.

1. Talk of conceptual content is ambiguous between *content-conceptualism* and *state-conceptualism* (on this, I follow Heck (2000)). State-conceptualism (SC) holds that the concepts *possessed by the subject* constrain perceptual content:

(SC) *S*'s mental state, *M*, at time *t* has *conceptual* content *p* iff at *t*, *S* possesses the concepts that characterize *p* and *S*'s possessing those concepts is involved in *S*'s being in *M*.⁵

Claims that what one can see depends on what concepts one has are not unfamiliar, and they have an echo in SC. Certain invocations of aspectual shape or seeing-as also suggest that concepts must be involved in perception. These are intuitively forceful ideas, yet we have not explicated these claims effectively because we do not understand what the involvement of concepts comes to.

For most philosophers in this area, the issue is whether *content*, understood as an abstract object, is conceptual and the condition at issue is content-conceptualism:

(CC) *S*'s mental state, *M* possesses *conceptual* content *p* iff *p* is constituted by concepts.

I shall, however, set CC aside. This might seem to cut my discussion loose of much of the philosophical debate, but this need not be a bad thing, especially if there is a substantive question about conceptual content, given SC, that has been not received the attention it deserves.

It has been possible thus far to refrain from specific claims about concepts, an actively debated topic in philosophy and psychology. In aiming for conclusions compatible with most extant theories of concepts, I aim to rely on commonly accepted premises. Thus, I begin with a truistic condition on concept possession: to have some concept *F* is to have a capacity to think about *F*s.⁶ This *thin conception* of concept possession, as we can call it, implies that there is a difference in cognitive capacities

⁵ Strictly speaking, we should speak of a conceptual *state* with content *p* rather than a state with conceptual content *p*. The first conjunct on the right is perhaps redundant.

⁶ This condition is accepted by both Fodor (2004) and Peacocke (2004).

between someone who has a concept, F, and someone who does not. Obviously, theories will differ on the precise nature of this capacity (this is what individuates them), but they are unified in a commitment to its presence.

Consider an individual who prior to some time, t , lacks the concept F and is thus unable to think about Fs, but who after some later time, t' , is able to think about Fs and thus has the concept. Can we provide a more informative description of what changes? First, the changes will often come about as a result of one's interaction with the environment, one's experience of the world over time whether in interaction with Fs or, more indirectly, in gaining information about Fs through other sources. These experiences leave their *trace* in the individual so as to put her in a position to be able to think about or otherwise react cognitively to Fs, traces that are typically stable. By speaking thus about stable traces in the subject's mind that reflect her experiences with Fs and ground abilities to react to Fs, we have specified, in the most general terms, a form of memory concerning Fs.

Our capacities to think about Fs then have as a necessary condition memory of this sort, what psychologists term *semantic memory*. Experimental work has shown that semantic memory, localized in the temporal lobe, is called upon in various cognitive tasks such as recognizing objects (F 's), categorizing them, understanding and constructing sentences about them, reasoning about them, and describing and otherwise depicting them.⁷ All of these tasks can plausibly be understood to depend on an ability to think about the objects in question, so a loss of that capacity for thought about Fs, due to deficiencies in semantic memory, should show itself in impairments in these abilities.

⁷ For a recent review, see Martin and Chao (2001).

There is a more complicated question of how semantic memory relates to concepts. Many psychologists take semantic memory, or the representations stored therein, to *be* concepts. In offering their theories (prototypes, exemplars, theory accounts, etc), psychologists aim to elucidate the structure of concepts. Indeed, many philosophers of concepts see the two approaches as directed at the same phenomenon.⁸ I favor a different view: psychologists, in being concerned with information processing in the brain and the relevant structures required for it, are delineating the neural grounds for concepts and conceptual thought. Accordingly, if we take (neural) information-processing to be below the level of the person (i.e. subpersonal), then semantic memory is the sub-personal correlate of concepts, concepts being, as philosophers hold, the constituents of thought—a personal level phenomenon if anything is. On this view, our capacity to actively think about Fs depends on semantic memory regarding Fs as at least a necessary condition. It is plausibly, however, to also take semantic memory as what realizes (or at least partly realizes) our capacities to think about Fs. This additional step is more controversial. I do not argue for it here, but treat it as a theoretical hypothesis.

We now have a way to link philosophical concerns about concepts with psychological theory. Philosophers, by their commitment to the thin-conception in characterizing thought, are describing a mental phenomenon that depends on a form of memory as a causal ground. Psychologists, in their investigation of concepts via semantic memory, are probing the structure of that ground. The upshot is that we can make use of empirical work on semantic memory to probe the involvement of concepts in various aspects of mind. If concepts are involved in certain mental processes, then

⁸ Notably, Jesse Prinz (2002). See the collection by Margolis and Lawrence (1999), especially their introduction for an overview.

affecting semantic memory should lead to alterations in those aspects of mind. This is an empirical consequence of the thin-conception of concepts and I shall consider work that is relevant to it below.

We must, however, consider what role concepts might play in vision (see section 3). Accordingly, it will be necessary to keep a firm grip on the personal-subpersonal distinction as applied to vision and attention. Although the distinction needs clarification, in this context I simply mean to signal a difference between conceptual processes of the person such as thinking versus information-processing of neural systems. When I speak of a personal level phenomenon, I shall explicitly mention the subject and speak of the subject's *seeing* and *visually attending*. When I speak of sub-personal phenomena, i.e. the operations of the cognitive or neural information-processing systems within subjects, I shall speak of *vision*, *visual processes* and *attention*. I shall also restrict the notion of *information* to neural processing. In these latter cases, the subject will not be mentioned.

The point is to understand how concepts are involved in vision, which in personal level terms, is to understand how a subject's seeing or visually attending to the world amounts to a state that counts as conceptual as per SC. But one way to get at this issue, given the link to the sub-personal notion of semantic memory, is to determine whether and how semantic memory affects visual and attentional information-processing. It is these sub-personal interactions between visual processing and semantic memory that form the causal ground of the involvement of concepts in a subject's visual experience, and this in turn will license the personal-level claim that when agents visually attend (for example) to a hammer, their seeing the hammer in that way is a conceptual state they are

in. That is the conclusion that will contradict those who say visual content is always nonconceptual, but can we get it?

2. In this section, I argue that vision must play a direct role in action in which case the input-only view is false. The picture to be opposed is where intentions function as the proximate causes of action so that perception, if it is to have an influence, affects action through intention as a causal intermediate. Yet intentions cannot do all the work that the picture calls upon them to do which is to guide and control action. Consider my intention to pick up the hammer on the table in order to fix a chair. On the Causal Theory, the intention causes, controls and guides the action of picking up the hammer, but how does that intention generate successful action? In the normal case and on the basis of that intention, we successfully pick up the hammer, we reach accurately for it and handle it in an appropriate way, yet, the content of the intention does not specify enough information that could plausibly suffice to control and guide the requisite finely-tuned bodily movements. For example, the intention does not encode the object's precise spatial location, shape, orientation relative to the body, and so forth. Such information is needed to guide and control action, or more specifically, the bodily movements required for action. It is mysterious, then, how intention leads to successful action given the limits of its content.

Why should we take there to be such limits on the contents of intention? Intentions have their place in our mental economy as products of practical reasoning. This is not to rule out intentions which just come to us (if there are such things), but the primary connection is to the agent's capacity for reasoning about what to do. What we typically reason about sets conditions on what we encode in our intentions. It is

implausible, however, that we always reason at such a fine level of detail about the trajectory of our bodily movements, the precise aperture of the grip needed to grasp an object, and the precise spatial location of things relevant to how we act. Yet proper action requires that information. If it is not encoded in intention, where does it enter?

Perhaps we should shift from a ‘descriptive’ intention, an intention to pick up *the* hammer, to a demonstrative intention to pick up *that* hammer. On the most plausible theories of demonstrative thought, having such thoughts concerning an object implies a capacity to locate it (see Evans, 1982, chp. 6; cf. Campbell, 2002). If so, the appeal to demonstratives in our intentions can help explain how the agent is able to locate the object in her action.

Yet, a demonstrative intention does not suffice on its own to explain how intentions lead to successful action. Consider a case where we tactilely demonstrate the hammer but otherwise have no information made available by vision (our eyes are closed). We may put our finger on the hammer’s handle and recognize the object as the hammer in question by its unique texture. On the basis of touching it, we can form an intention to pick up *that* hammer. Yet in placing our finger on the hammer, we do not gain sufficient information as to other properties that we must respond to in order to act effectively on it, even as touching allows us to spatially locate it. We lack information as to the orientation of the hammer’s head or the location of the end of the handle. Of course, we can get this information tactilely by actively exploring the object but this further exploration is not needed to establish the demonstrative link.

The tactile case demonstrates that an appeal to the mere possibility of a demonstrative intention does not suffice to explain how we act successfully on objects.

We need access to more information than just the object's location. In our case, this information is usually made available in a glance when we visually take in the hammer. By looking at it, we register its action-relevant properties such as shape, location, size, orientation and so on. This information can then be called upon to guide and control one's reaching for the hammer. It is important, however, to recognize that the information is not encoded in one's intention, for we do not reason about it. What we intend is merely to pick up *that* hammer in order to use it; our intentions need not encode anything more than what we have decided to do on the basis of reasoning.

Moreover, why should we take the further step to encode all this information in our intentions when it is already available just by maintaining a visual link with the object? If intentions do have a controlling role to play with respect to action, they do so not by encoding all the relevant information needed to generate appropriate movements, but rather by bringing the relevant visual information, readily available, directly to bear in the production of bodily movement. Intentions, on this view, function as a conduit linking visual information to the control and guidance of bodily movements. They thereby causally *mediate* but are not causal *intermediaries* between vision and action. Perception is not one step away from action, with intention in between; rather, perception directly links to action, though this is enabled by intention.

By recognizing that vision must play a direct role in action we provide a framework to answer (3): perception is brought to bear by intentions to directly control bodily movements. This now raises questions about the character of the agent's vision in action to which I now turn.

3. In this section, I shall address (2) and outline how conceptual capacities might be needed to modulate visual processing in a way that explains how concepts can be involved in vision. The basic ideas concern what I shall call the *focus* and *parsing* functions of conceptual capacities which constitute a way of visually attending to the world. I then review empirical evidence that these conceptual functions are exemplified in vision. The upshot is that agents visually attend to an artifact *as the artifact that it is* in a way that involves the concepts they possess (e.g., they visually attend to the hammer as a hammer). Hence, they are in a conceptual visual state, as per SC.

Earlier, I argued that vision is given direct control in action. What we see, however, must be organized in order to serve our goals and ensure the successful execution of action. After all, in a given glance we see a variety of objects only a subset of which is relevant to our current goals; without some way of focusing on the relevant object, we would be unable to act. I shall argue that concepts can play an organizational role at this stage. Returning to our earlier case, we may need to use a hammer to fix the chair, but what we see is a work-table littered with objects. So there is an initial problem: our attending to the right object. How is this done?

Given our goals, certain objects become salient to us. Thus, if one intends to use a hammer, the hammer comes to stand out as opposed to a screwdriver. Note that here, a demonstrative intention is not needed; one intends to use a hammer in a certain way, and as a result, we come to visually attend to that hammer in a way that guides the production of appropriate bodily movements without modification of our intentions. The simplest explanation of this is that there is a content-specific relation between one's intention and one's attending to the hammer: being in a state of intention to use a *hammer* leads one to

attend to it. Since intentions are conceptual-states, the concept HAMMER deployed in our intention can play a role in organizing what we see in a way that enables us to attend specifically to the hammer. One mechanism for this is that the appropriate aspects of semantic memory, correlated with our concept HAMMER, organizes the plethora of information that is processed by the visual system so as to underscore information regarding the hammer against information relating to other objects, making the former available for further processing.⁹ This interaction between the visual system and semantic memory is the sub-personal grounds of our attending visually to that hammer. Concepts have a *focusing* role to play by enabling the selection of the relevant object against irrelevant ones.

Still, if our aim is not only to locate a hammer but to *use* it, then we must also visually grasp the hammer's functionality, and this requires *parsing* its various properties. This echoes an earlier point that a demonstrative intention could not on its own lead to successful action as that would only guarantee our spatially localizing the object. More is needed to ensure effective hammering: the various seen properties of the hammer must be organized in the right way. This poses an analogous problem to the earlier problem of attending to the right object amidst a clutter of irrelevant objects. We must also attend to the hammer's relevant functional properties against a clutter of irrelevant properties. Our concept HAMMER can provide resources to produce an accurate functional parsing. A similar mechanism as deployed in focusing may be in play: the involvement of relevant aspects of semantic memory in visual processing can lead to the selection not only of

⁹ The ideas spelled out here echo the current dominant account of perceptual attention, the biased-competition model (Desimone and Duncan, 1995). On this view, we attend to those objects whose information processing streams are biased in a way that allows them to gain access to limited processing resources.

information concerning hammers as opposed to other objects, but in this case, the selection of the relevant properties of hammers and their relations as opposed to irrelevant properties of the same object. At the personal level, drawing on our concept HAMMER, the link between the location of the handle and the orientation of the hammer head, relative to the agent's body is made salient. Here, a *parsing* role in visual processing is attributed to concepts, or their subpersonal correlates, and the result is that the agent comes to attend visually to the object in a specific way.

My brief description of focusing and parsing will elicit questions about how the correlated aspects of semantic memory can organize visual information. That is a mechanistic question, however, for vision science. What I am relying on here is the possibility that having a concept enables us to deal adequately with the objects that fall within its extension, whether this means drawing the right inferences concerning it or reacting appropriately to it. On the practical side, reacting appropriately to an object in action involves seeing it for what it is just as, on the theoretical side, drawing the right inferences about an object involving understanding it for what it is.

Given a context for acting, there are then two postulated roles for concepts in vision: an object must be localized and it must be parsed visually. If concepts do play a role in organizing how we see, then this seems to fit the *involvement* requirement of SC on a reasonable reading of 'involvement.' Specifically, involvement comes to getting the agent to attend to the object *as the functional object it is*. Thus, in our test case, we might say that to act effectively on the hammer, the agent must see the hammer *as a hammer* where this sort of aspectual seeing amounts to a conceptual visual state in the sense of SC. In other words, we can give a specific sense to the idea of seeing something *as a*

hammer. This implies that the subject has the concept HAMMER and that her possessing the concept is involved in that state.

In my discussion of focusing, I explicitly mentioned intentions directed at certain targets which result in attending to those targets, and this might prompt the thought that the involvement of concepts here is in intention rather than in vision. But there is another way of describing the situation, namely that what intentions achieve is to activate the relevant concepts (and their correlates in semantic memory), and in doing so, intentions make these concepts available to play a role in vision which I argued must be directly involved in producing action. Accordingly, both intention and vision will draw upon the same conceptual capacities, even if this is achieved by first deploying the relevant concept in intention. The further involvement of concepts in vision is especially apparent in *parsing*, which depends on a dynamic and direct interplay between the visual system and semantic memory. After all, one must not only locate the hammer, but also locate the hammer's handle and grasp it in a way that is sensitive to the tool's orientation relative to its target and to the agent's body, information that is present in the content of visual processing but must be brought to bear in an organized way to deal with a specific object (the concept HAMMER must be applied to a specific hammer visually presented).

I have suggested that concepts are involved in vision by reflecting on what is required for acting effectively, given the complexity of what we visually take in, and I have explained what such involvement comes to. One might wonder, however, why concepts need be involved, why the object must be recognized conceptually so as to allow direct visual control of action. After all, having processed the object's form visually where this is understood to yield a structural but not conceptual grasp of the

object, we can then rely on mechanisms established by habits and practice (skills) to link to bodily movement with no detour through concepts. We simply pick up on the object's *affordances* which, as J. J. Gibson said, "is not to classify an object" (1979, p. 134). We have, then, two opposing accounts. Which is correct? I have argued, however, that if concepts play a role, then probing semantic memory ought to reveal this. I shall briefly recount two experimental results that support the involvement of concepts in vision.

The first comes from patients with semantic dementia, i.e. damage to their semantic memory system which incapacitates them with respect to their conceptual abilities. Although it has been claimed that these patients operate normally with objects for which their conceptual abilities are degraded, careful testing reveals that they are impaired in appropriate use of objects in a way that tracks their conceptual impairments.

In a study by John Hodges and colleagues, 9 patients with varying forms of semantic dementia were tested for their ability to use everyday objects (2000). The subjects were assessed for abilities to understand the use of objects and were shown to suffer varying degree of conceptual impairments regarding those objects. Subjects were then tested for their abilities to demonstrate the use of the objects (including correct grasp, orientation and movement). The tests showed that failure to use objects directly correlated with semantic disabilities, suggesting that semantic memory was required for object use. Thus, if one demonstrated a deficiency in semantic memory with regard to scissors, one would grip scissors incorrectly or mime the wrong movements with them.

One may worry that with degraded conceptual capacities come degraded thoughts, specifically an inability to form intentions. But to attempt the tasks, the patients had to understand the instructions (e.g., show us how to move this object), so

presumably had to be able to intend to do just that. Since these individuals were presumably able to attend to the objects in question, the defect would seem to reside in parsing. Thus, it seems that even if one is able to form the requisite intentions (e.g., to demonstrate the use of that object presented to me), the absence of conceptual knowledge devastates appropriate action by affecting visual parsing.

Sarah Creem and Dennis Proffitt (2001), testing normal subjects, showed that semantic memory is required for functionally appropriate action. They conducted experiments where subjects were required to grab objects by the handles which were placed before them during performance of a semantic recall task. The results are striking. If subjects simultaneously perform a memory task (recalling the matched word to previously memorized word-pairs on being given one of the two words), they failed to grasp the object appropriately. The same results were obtained if the subject's were asked to produce words in response to prompted words by free-association. However, when the distractor task was a non-semantic, spatial imagery task, objects were used appropriately.

To rule out the possibility that the semantic tasks were somehow more difficult than the spatial imagery task, a second experiment was conducted where the task was a visuo-spatial task, namely to track a moving target on a computer screen by moving a cursor with a mouse. The distractor tasks were the same as in the first experiment, and the opposite affect was found in this case. The semantic memory tasks had no effect on the successful tracking of the target, but the spatial imagery task did negatively affect the subject's ability to track. Given this double-dissociation, it seems that semantic processing must inform the visual spatial processing in the service of producing bodily

movements so as to use objects appropriately. Semantic memory is necessary for the visual processes at issue. Given my theoretical hypothesis that semantic memory provides at least part of the causal basis for the operation of concepts, we can conclude that concepts are involved in our seeing the world as we act. Thus, I claim that we have provisional evidence that concepts are involved in vision in a way consistent with SC and that the relevant visual states are conceptual-states. Vision, then, is at least sometimes conceptual.

To conclude, in respect of our opening questions I have argued that perception must be directly involved in action, that concepts are involved in vision in action on artifacts and that conceptual involvement is rooted in the focusing and parsing roles that concepts can play vis-à-vis vision, roles that lead to a form of visual attention.¹⁰ We reached these results by exploring the link between action and perception. Strikingly, the philosophy of action has implications for the philosophy of perception.

¹⁰ I thank Marlene Behrmann, Liz Camp, John Campbell, James Genome, Hemdat Lerman, Edouard Machery, Linda Moya, David Plaut, David Ralph, Sarah Schomstein, John Searle, Helen Steward, Jay Wallace and an anonymous reviewer for the 9th Oxford Graduate Student philosophy conference for their responses to previous versions of the ideas presented here.

Bibliography

- Byrne, A. (2004). Perception and Conceptual Content. Contemporary Debates in Epistemology. E. Sosa and M. Steup. Oxford, Blackwell Publishers Inc.
- Campbell, J. (2002). Reference and Consciousness. Oxford, Oxford University Press.
- Creem, S. H. and D. R. Proffitt (2001). "Grasping Objects by their Handles: A Necessary Interaction Between Cognition and Action." Journal of Experimental Psychology: Human Perception and Performance 27(1): 218-228.
- Cussins, A. (1990). The Connectionist Construction of Concepts. The Philosophy of Artificial Intelligence. M. Boden. Oxford, Oxford University Press: 368-440.
- Desmione, R., and J. Duncan (1995). "Neural Mechanisms of Selective Visual Attention." *Annual Review of Neuroscience* 18, no. 193-222.
- Evans, G. (1982). The Varieties of Reference. Oxford, Oxford University Press.
- Fodor, J. (2004). "Having Concepts: A Brief Refutation of the Twentieth-Century." *Mind and Language* 19, no. 1: 29-47.
- Heck Jr., R. G. (2000). "Nonconceptual Content and the "Space of Reasons". " Philosophical Review 109(4): 483-523.
- Hodges, J. R., S. Bozeat, et al. (2000). "The Role of Conceptual Knowledge in Object Use: Evidence from Semantic Dementia." Brain 123: 1913-1925.
- Hurley, S. (1998). Content and Environment: Parallels between Perception and Action. Consciousness in Action. Cambridge, MA, Harvard University Press: 245-284.
- Margolis, E., and S. Laurence (1999). "Concepts and Cognitive Science." In *Concepts: Core Readings*, edited by Eric Margolis and Stephen Laurence, 3-81. Cambridge, MA: MIT Press.
- Martin, A., and L. L. Chao. (2001). "Semantic Memory and Brain: Structure and Processes." *Current Opinion in Neurobiology* 11: 194-201.
- Noë, A. (2004). Perception in Action. Cambridge, MA, MIT Press.
- Peacocke, Christopher. "Interrelations: Concepts, Knowledge, Reference and Structure." *Mind and Language* 19, no. 1 (2004): 85-98.
- Prinz, J. (2002) *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.

Visual Attention, Conceptual Content and Doing it Right

Stalnaker, R. C. (2003). What might Nonconceptual Content Be? Essays on Nonconceptual Content. Y. H. Gunther. Cambridge, MA, MIT Press.

Wu, W. (ms). "Making Sense of Attention, Action and Agency in the Causal Order."