

RESPONSE TO REPLIES

Reconstruction and Purpose

MICHAEL J. TARR*

Department of Psychology, Yale University

AND

MICHAEL J. BLACK

Department of Computer Science, University of Toronto

Received September 1, 1993; accepted November 2, 1993

1. ALL VISION IS GOOD VISION

We greatly appreciate the time and thought the authors of the replies have put into their commentaries on our paper [1–9]. We are concerned, however, that there has been some misinterpretation of our goals. We emphasize that we are attempting to represent a *conservative* position on the study of vision. It is not that we believe that current and past paradigms are ideal and should be maintained to the exclusion of all other approaches, but rather that we feel that the goal of reconstruction and scene recovery is viable and should not be completely abandoned. We claim this with full knowledge of the fact that these and related terms, most notably “general vision,” cannot be defined precisely at present (despite the best efforts of [3, 9])—a fact that advocates of the purposive paradigm have used as a justification for a radical departure from the current paradigm. Perhaps our concern in this regard is unwarranted, thereby rendering our caution meaningless. Yet even within the commentaries it is clear that there exist strongly contrasting points of view. Broadly characterized, there are three apparent stances:

- We do not know what general vision is, therefore we should forego studying it.
- We do not know what general vision is, therefore we should study it exclusively.
- We do not know what general vision is, but all vision is good vision.

What we wonder after reading these commentaries is whether anyone seriously believes the first or second positions. Moreover, no matter which position one adheres to, everyone seems to agree that we have no idea of what

general vision really is. Yet *something* motivated us to rush to its defense. What prompted our original concerns were exclusionary statements at the expense of the reconstructionist approach, including phrases such as “radical change,” “paradigm shift” [10], “the brain is nothing but a set of behaviors” [2], “general-purpose vision is a chimera” [2, 3], and “this traditional reconstructionist approach is questioned and the paradigm of purposive–qualitative vision is offered as an alternative” [11]. Indeed, some of our reaction may stem from a confusion in the field itself over whether there is an incompatibility between the old and the new. For instance, Brooks¹ writes in the forward to a recent book titled *Active Vision* [10] (for the moment putting aside the question of whether purposive vision and active vision are equivalent) the following passage:

To be sure, the papers in this book do not represent a complete break with traditional work in computer vision. There are clear continuities with the work of the last twenty years. Three dimensional models and symbolic descriptions still play a role in many of the papers. The book does, however, represent a change in emphasis; vision is no longer a passive process, instead it interacts with the dynamics of the world, becoming an active participant. (p. xi)

On the other hand, a mere six pages further into the book, the editors state that active vision “. . . represents a paradigm shift (Kuhn, 1962), a radical change of emphasis on what is considered important in vision” [10, p. xvii].

Due to such strong claims, as well as apparent inconsistencies, we were motivated to state what Brooks put so succinctly, that elements of the purposive approach are not that different from what we have been doing all along and, moreover, that the two approaches are not competi-

* Please address all correspondence to Michael J. Tarr, P.O. Box 208205, New Haven, CT 06520-8205. E-mail address: tarr@cs.yale.edu.

¹ The very same theorist who wrote articles entitled “Elephants Don’t Play Chess” [12] and “Intelligence without Representation” [13].

tors, but rather offer compatible perspectives on the study of vision. Indeed, in contrast to previous work, the commentaries exhibit a far less strident tone and much more accommodation of differing viewpoints. What appears to be happening is that upon self-analysis neither the purposivists nor the reconstructionists can maintain an extreme stance (ourselves included); rather, introspection reveals a somewhat more moderate position in which elements of both approaches emerge. In retrospect, this is not entirely surprising, for as Jain's PIPO principle suggests, all vision systems should have *some* goal. Likewise, all vision systems must reconstruct *some* properties of the scene. Thus, purposivists must acknowledge that reconstruction will occur and reconstructionists must acknowledge that there is a purpose to vision. In essence, what may lie at the heart of this discussion is the question of exactly what the goal of vision is. This is a question that we cannot generally answer—what we can do is reiterate our point that the goal, sometimes implicit, of many vision researchers is matching the visual performance of humans. This is clearly evident in that advocates of the purposivist approach [11], as well as many archetypical reconstructionists (e.g., Marr [14]), motivate their work by comparing the capabilities of current computer vision systems to those of humans and other animals. Yet this certainly does not mean that this is the *only* goal available to vision researchers. Indeed, our statement that the purposive paradigm is “better suited for understanding and mimicking the overall visual behavior of frogs rather than humans” is not meant to belittle the approach at all (and it is unfortunate if it was taken this way). Rather, it is meant to raise the point that, in our estimation, an extreme version of purposivism, e.g., the purposive paradigm, will not yield answers that will lead to the complete understanding of *human* vision.

1.1. *The Elusiveness of General Vision*

One simple definition of general vision is vision as performed by humans. As pointed out by Aloimonos, Jain, Sandini and Grosso, and Tsotsos, this obscures the point that human vision is not truly general vision but, like the vision of bees and frogs, has evolved to satisfy a purpose. It is appealing to suggest, of course, that human vision is *close* to general vision. But this is an assumption that cannot be proved and may already have been disproved [15]. Likewise, general vision (or human vision) cannot possibly be implied to be the literal reconstruction of the scene—as clearly shown by Aggarwal and Martin, a perceptual system (or a program) is unable to account for all observable properties or purposes within a recovered model.

While we agree that general vision is not a well-defined concept and, in particular, may be impossible to define as a competence theory, we suggest a strategy of avoid-

ance—sidestepping the issue completely. In this regard, we agree with Brown that a *performance theory* of general vision exists and is in fact defined by the vision we have (illusions do not render human vision purposive—they simply indicate that visual perception involves statistical inferences about the scene that are non-accidental, but that are violated by devious experimenters). Note that this definition leaves open many possibilities because, as pointed out by Aloimonos, our knowledge of human vision is extremely limited.² However, as mentioned, this does not mean that human vision itself is purposeless or that there are not properties of the scene that we fail to reconstruct or recover under a given set of conditions. Rather, as discussed in Section 2, it implies that human vision is structured in a manner that allows the selection of a variety of mechanisms that may be assembled to solve novel, and heretofore unexpected, visual tasks. Thus, while general vision is not well understood, the claim that it may be operationalized as human vision leads to two conclusions:

1. General vision is elusive, in part, because we do not yet have an adequate understanding of human vision and because human vision is a highly complex and adapted system.

2. Human vision offers a model on which to base and against which we can measure machine vision. As a consequence, assessments of the performance of machine vision systems must address not only the question, “Does it perform the task as well as human vision?”, but also “Can it learn to perform unexpected tasks as well as human vision?”.³

1.2. *The Elusiveness of Purposive Vision*

While everyone seems to agree that general vision is an elusive concept, the majority of the commentators seem to be comfortable with the concept of purposive vision. In the broadest sense, we agree that vision is purposive; however, we are specifically unsure of the purposive paradigm in its most extreme behaviorist incarnation. Like Brown, we wonder whether purposive vision is composed of more than one thing. On the one hand, because the purposive paradigm seems to require the declaration of very specific and narrow goals, it appears to exclude more general purposes, for instance, recovering the shape of an object for possible grasping, recognition, or reasoning. On the other hand, declaring the goal of a system to be shape recovery is still purposive; it is simply that the purpose encompasses a wider range of tasks for

² Indeed, one issue this raises is that perhaps computer vision would profit from learning more about what the human visual system *can* do.

³ Note that this follows Jain's stricture that we ask what a system can do, rather than when it will fail—here, all we have added is, “What else can it do?”

which knowledge of shape is required. At this point, a purposivist might argue that once such tasks have been identified, each should be solved “directly,” using the minimal information required [11]. The more general module is then dispensed with, replaced by an array of qualitative mechanisms that suffice only for their single, specific purposes.

While this may seem straightforward, many questions arise. First, is such an approach likely to provide mechanisms that are truly different from the more general case? More to the point, what does it mean to solve such problems “directly?”⁴ While it is possible that there are instances in which a simple invariant specifies the information sufficient for task completion, and therefore may be used “directly,” there are likely to be instances where no such invariants exist (for example, variation of features with viewpoint [18]). In such cases, some recovery must occur no matter what the task; consequently, a great deal of redundancy may arise between purposive mechanisms, each disregarding different commonly recovered properties of the scene. A more parsimonious solution may be to instantiate a single shared representation that is then used selectively according to the given task.

A second question that arises is how assemblies of inflexible qualitative mechanisms can accommodate new, unexpected visual tasks? As pointed out by Christensen and Madsen [4], “different techniques are utilized in an efficient manner and not merely due to simple synthesis.” For instance, if shape is used only insofar as is necessary to find a good grasp, how could this mechanism be combined with others to support a task requiring different properties of the shape (e.g., object recognition)? One answer may be that this is a new and distinct purpose, therefore a new and distinct task-specific mechanism should be implemented. However, as discussed in Section 2, purposive vision cannot continually invent mechanisms for every possible task unless (a) recovery mechanisms have already made relevant properties available or (b)

every possible visual goal has been anticipated and instantiated as a distinct mechanism.

Much as we argued that the ideal of general vision is unattainable, we suggest that ideal purposive vision is equally unattainable. Similarly, while human vision provides a simulacrum of general vision, other biological examples provide a metaphor for purposive vision. Here we suggest again that purposive vision may be typified well by biologically motivated task-specific mechanisms. Several characteristics make this so: first, as pointed out by Tsotsos, such mechanisms are predicated on the need for speed—the narrower the purpose, the faster the task can be performed; second, such mechanisms are inherently qualitative—they generally measure the presence or the absence of a single property; third, such mechanisms disregard most of the information available in the optic array—they use only the information that is relevant to the defined purpose. However, even for such systems the possibility that they are both purposeful and reconstructive may be worth considering. For instance, even the bug detectors within the frog retina both serve a specific goal and provide a representation of elements of the scene (albeit a highly constrained representation). For while their purpose may be to detect bugs, their implementation requires recovering, representing, and possibly responding to instances of moving dark spots—the fact that nothing else may be done with this encoding does not in any way diminish the fact that this is still a representation of some property of the scene. Of course, all this example may establish is that even the purposive paradigm entails recovery (a point that was never really subject to debate). However, we claim that this example also serves to illustrate that recovery leads to reconstruction: perceptual systems can only represent properties (e.g., detected moving dark spots) and not actual physical objects (e.g., bugs). As such, the representation is a reconstruction of the physical world, not simply a recovered property or an invariant in the world.

Finally, these claims, in particular that purposive vision may be operationalized as highly task-specific vision, lead to two conclusions:

1. Purposive vision is elusive, in part, because we do not yet have an adequate understanding of when task-specific, time-critical processing should be implemented and because we do not understand the computations that may arise from the combination of many such processes.

2. Task-specific biological vision can offer another model on which to base and measure machine vision. Assessments of the performance of machine vision systems must address questions such as “Does it perform the task as well as specialized systems?”, “Can it do so in a time-efficient, robust manner?”, and “How will it interact with other such systems?”.

⁴ The Gibsonian [16] notion of “direct perception” will not help us here, for it eschews *all* computation, including even the task-specific sort implemented in the purposive paradigm. To a true Gibsonian, all visual tasks may be solved simply by picking up invariants inherent in the optic array. Moreover, even if one allows for *some* computation, the claim that task-specific “features” may be directly perceived (as suggested by Edelman) is misleading, since the very idea of a feature indicates a great deal of inferential processing—particularly if one is talking about features at the level of cortical brain regions (MT) that necessarily receive visual input only after extensive processing and recovery has already occurred in the retina, LGN, and other intermediate structures. Even if one shifts the example to highly task-specific mechanisms within the earliest layers of vision, for instance, retinal bug detectors, the “pick up” of the pertinent features will still rely on computation. Likewise, the Gibsonian idea that objects inherently “afford” their functions is ill-defined and meaningless—we can readily learn to identify and use arbitrary objects that have been assigned arbitrary functions (see Ullman [17]).

1.3. *What Purposive and General Vision Are Not*

While there is some agreement that purposive and, certainly, general vision are elusive concepts, there is less of a consensus on what these approaches *are not*. Many of the commentaries attempt to associate one or another of several different computational dichotomies with that of purposive/general vision. However, it is our contention that such associations are at best stereotypes and at worst misleading. Specifically, at one time or another we have seen purposive vision paired with characteristics such as active, top-down, and distributed, while general vision has been paired with characteristics such as passive, data-driven, and symbolic. Of course, at a superficial level, these pairings do capture the history of the two schools. Why then do we refuse to accept these characteristics as defining properties? First, if we examine the active versus passive distinction, it is clear that there is nothing inherent in the goal of reconstruction that precludes an active perceiver. Indeed, as pointed out by Brown, one of the most significant contributions of the purposive paradigm has been to raise consciousness about active vision throughout the community. But while active vision has a purpose in that guidance is presumably motivated by a goal, it is not purposeful in the sense of being restricted to seeking information for only a single, specific task; rather, information may be relevant to many aspects of vision, including, for example, the recovery of scene structure. This brings us to the second distinction, that of top-down versus data-driven processing. It would of course be appealing if vision could function successfully in a purely bottom-up fashion—much as Marr [14] proposed in his reconstructionist account of vision. However, many problems in human vision are helped by the introduction of context; for instance, it has been demonstrated many times that object perception is facilitated by consistency within a scene, e.g., knowing you are observing a restaurant. Additionally, it is clear that visual attention is a limited-capacity resource allocated by humans according to top-down constraints in an active manner. But neither of these phenomena imply strict purposive vision—scene constraints may, for instance, simply reduce the model base for indexing, and attention may be allocated for general tasks, such as taking in an entire vista. On the other hand, purely bottom-up processing may sometimes be purposeful. For instance, a highly task-specific process may instantiate specific inferences learned over the evolutionary history of the organism; thus, the process may appear to have top-down constraints, but is in fact purely data-driven. Finally, it has been suggested that the goal of reconstruction is best understood as the derivation of symbolic description, whereas the purposive paradigm is most often associated with distributed representations. However, this distinction is a false dichotomy in that both

computational architectures can accommodate either type of approach. Indeed, there are many instances where neural nets have been used to solve reconstruction problems. Thus, as we suggested in our paper, the reconstructionist approach is compatible with many elements typically associated with purposivism, for instance, active top-down processing and distributed representations.

2. NATIVIST VS ON-LINE PURPOSIVISM

We now turn briefly to the can of worms known as representation. We admit, as pointed out by Fischler, that we have said little about the nature of the representations we have advocated. Indeed, as mentioned above, because we take human vision to be an example of general vision and because we currently do not know a great deal about human vision, we cannot offer a well-specified theory of representation. We do, however, believe that we can offer two perspectives on the possible architecture of perceptual systems. Specifically, given that vision is used to accomplish certain purposes, these can be achieved through either *nativist* or *on-line* purposivism. The nativist perspective is consistent with the way we have characterized the purposive paradigm to this point: visual tasks are accomplished by an array of specialized narrowly defined task-specific mechanisms. The primary point is that each mechanism has evolved to meet a single purpose and, as such, operates in a predetermined and hard-wired manner. In contrast, the on-line perspective is characterized by visual mechanisms that are not narrowly defined; rather, there exists an array of mechanisms that are capable of recovering properties of the scene regardless of purpose. Relevant information from such mechanisms is then combined on-line in a flexible, yet task-specific, manner. The primary point is that on-line purposivism requires the availability of appropriate information not just in the optic array, but in the representations of the recovered optic array. Furthermore, unlike theorists who have no intention of modeling, simulating, or otherwise implementing such systems, we cannot appeal to some neo-Gibsonian mysticism in which the environment affords function or the requisite information is “directly perceived”—information must be perceived in order to be used. The distinction between nativist and on-line purposivism is important for understanding the difference between the purposive paradigm and general vision: the former is a nativist account of how goals are realized in vision, while the latter is an on-line account. To reiterate, all vision, including general vision, has a purpose, but in the case of general vision goals are not predetermined; therefore, they must be handled by more inclusive mechanisms that may be combined to arrive at task-specific solutions. Indeed, it is our belief that this view of vision

is consistent with inferences drawn from evolution⁵ and with theories of human visual processing.

2.1. *Visual Cognition*

Aloimonos points out that humans have a rich set of cognitive capabilities that appear to produce “spectacular, amazing, and general” performance. Moreover, as he suggests, because our visual systems interact with such capabilities it is difficult to determine which parts of performance should be attributed to these wonderful skills and which parts should be attributed to vision. What he seems to suggest is that vision is actually a bit player in this drama—that our vision is actually quite similar to the vision of bees or flies and does not play a central role in our apparent general cognitive skills. Indeed, for many years this was precisely the stance held by most psychologists: vision as an input system. In both instances, there appears to be a willingness to confer a unique status to human linguistic and reasoning skills, but not to human perceptual systems. However, since the late 1960s there has been a marked change in this thinking—most psychologists now hold that many of our cognitive capacities are perceptually specific. In particular, advances in the study of mental imagery have led to the evolution of a new subdiscipline known as “visual cognition” [19]. Crucially, the major tenet of this area is that our long-term visual representations retain many properties isomorphic to their physical manifestations [20], and, furthermore, that such representations may be used for many unanticipated cognitive tasks, including recognition, spatial reasoning, and visual problem solving [21]. For example, one property of visual images is that they can be *reinspected* in order to derive new explicit properties heretofore only implicit in the representation (e.g., whether a German shepherd has floppy or pointy ears) [22]. Thus, from a nativist perspective it is difficult to accommodate evidence that humans retain representations of objects that may be inspected, manipulated, and used in visual problem solving for unanticipated and novel purposes.

⁵ Not surprisingly, advocates of both the reconstructionist and purposive approaches cite evolutionary analyses as evidence in their corner. In some sense, both sides *must be* correct—because adaptation has endowed different species with different structures for solving similar problems depending on context and evolutionary history, there is no single biological “solution” to the vision problem. Claiming otherwise is akin to suggesting that evolution is consistent with only one side of the herbivore/carnivore debate for solving the problem of what to eat. Thus, there are examples of both narrow single-goal solutions and general multiple-goal solutions. Of course, as with all evolution-based arguments, our argument for why human evolution indicates a general solution to vision comes down to a “just so” story.

2.2. *What We Do Not Need*

Another characteristic of human cognition is that empirical results suggest a great deal of what we initially perceive is apparently retained in memory [23]. These findings have been codified into a theory known as *proceduralism* that suggests our memories are highly specific to the contexts in which they were originally experienced [23]; given the appropriate task, numerous seemingly irrelevant properties of the scene may be shown to influence later processing and, therefore, are retained. Such findings are clearly incompatible with the purposive paradigm—specifically, the nativist perspective suggests that even the most flexible and sophisticated representations are highly correlated with particular behaviors and, as such, are not general purpose, but follow the principle of “most commitment”—that is, “throw out what you do not need” [2]. In contrast, the procedural theory of human memory follows the principle of “least commitment”—retain everything. Obviously, such a theory poses significant problems of complexity; it remains to be seen how computational and neural models of perception and memory will account for the fact that empirical tests continue to expand the range of visual properties that are not discarded. Nevertheless, as a viable model of human memory with empirical support, proceduralism raises serious concerns about any claims that human vision is composed solely of predefined purposive elements that discard properties of the scene irrelevant to their given task.

3. TRUE CONFESSIONS

We end this reply on perhaps an even more agreeable note than the call for cooperation that ended our original paper. We feel that the current and, hopefully, future, discussions are healthy for the study of both biological and machine vision. In our own work, it has prompted us to reflect more upon the specific purposes of the aspects of vision we study and, moreover, to consider these purposes in the context of an active, dynamic perceiver. For instance, the first author’s work focuses on human object acquisition, representation, and recognition. Work in this area has traditionally addressed these issues in terms of static views [24, 25]. However, in reality, the human perceiver encounters objects in an active, constantly changing environment over both space and time. Furthermore, interaction with objects is not passive but involves active exploration based on perceived properties of the scene, current context and task, and prior knowledge. To this end, new psychophysical paradigms are being developed in which subjects manipulate and explore three-dimensional objects in a dynamic and interactive manner. Issues include the difference between passive and guided explo-

ration; the contextual factors that guide exploration; and the role of task and object function in structuring representations. These ideas have been shaped in part by trends in the computer vision community toward purposive and active approaches⁶ that, interestingly, were themselves shaped by theories of biological vision [16].

Similarly, the second author's work on motion estimation is providing the foundation for a number of experiments in active perception. For example, a computational model of early boundary detection from motion is being developed that allows an active robot to use its own motion for the detection, disambiguation, and tracking of object boundaries. In this research, active recovery and reconstruction work together and complement each other; egomotion permits the detection of object boundaries and the recovered surface boundaries support actions such as obstacle avoidance and grasping. In addition, research on the estimation of multiple motions is being applied to problems in active perception. Consider a forest-management robot that has the task of locating healthy saplings that are in danger of being choked by surrounding growth. While a single image may present a jumble of indistinguishable branches occluding each other, by using self-motion the robot can segment a scene into distinct layers corresponding to saplings at different depths. These examples illustrate how in our own research the reconstructive paradigm can naturally support tasks in active vision and how an active/purposive approach can aid in the recovery of scene properties.

Finally, why should we, as Brown states, "allow the nose of a goal-driven, purposive camel into our tent?" It is because, in spirit, purposive vision, rather than the purposive paradigm, allows for a broader definition of a goal—the recovery of some generally useful physical property of a scene, for example—and we have always supported the idea of studying vision in terms of some goals. Thus, perhaps the best answer is that this was our intention from the outset; the fact is that we have a very large tent and the purposive camels have been with us all along. Under this "big top" we want both purposivist and reconstructionist camels to feel at home; they can both put on a pretty good show.

ACKNOWLEDGMENTS

The first author was supported by a grant from the Air Force Office of Scientific Research, Contract F49620-91-J-0169. The second author

⁶ As pointed out by Brooks [10], these new approaches have been made possible by advances in computer hardware; likewise, new psychophysical paradigms have become feasible due to advances in computer technology.

was supported by the Natural Sciences and Engineering Research Council of Canada. We thank Laurie M. Heller for her helpful comments and advice.

REFERENCES

1. J. K. Aggarwal and W. N. Martin, The role of R & R in vision: Is it a matter of definition? *CVGIP: Image Understanding*.
2. Y. Aloimonos, What I have learned, *CVGIP: Image Understanding*.
3. C. M. Brown, Toward general vision, *CVGIP: Image Understanding*.
4. H. I. Christensen and C. B. Madsen, Purposive reconstruction, *CVGIP: Image Understanding*.
5. S. Edelman, Representation without reconstruction, *CVGIP: Image Understanding*.
6. M. A. Fischler, The modeling and representation of visual information, *CVGIP: Image Understanding*.
7. R. Jain, Expansive vision, *CVGIP: Image Understanding*.
8. G. Sandini and E. Grosso, Why purposive vision?, *CVGIP: Image Understanding*.
9. J. K. Tsotsos, There is no one way to look at vision, *CVGIP: Image Understanding*.
10. A. Blake and A. Yuille (Eds.), *Active Vision*, MIT Press, Cambridge, MA, 1992.
11. J. Aloimonos, Purposive and qualitative active vision, in *Proceedings, 10th International Conference on Pattern Recognition, Atlantic City, NJ, June 1990*. Vol. 1, pp. 346–360.
12. R. A. Brooks, Elephants don't play chess, in *Robotics and Autonomous Systems*, Vol. 6, pp. 3–15, Elsevier, Amsterdam, 1990.
13. R. A. Brooks, Intelligence without representation, *Artif. Intell.* **47**, 1991, 139–159.
14. D. Marr, *Vision*, Freeman, New York, 1982.
15. J. K. Tsotsos, Analyzing vision at the complexity level, *Behav. and Brain Sci.* **13**, 1990, 423–469.
16. J. J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, 1979.
17. S. Ullman, Against direct perception, *Behav. and Brain Sci.* **3**, 1980, 373–415.
18. J. B. Burns, R. W. Weiss, and E. M. Riseman, View variation of point-set and line-segment features, *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(1), 1993, 51–68.
19. S. Pinker, Visual cognition: An introduction, *Cognition* **18**, 1984, 1–63.
20. R. N. Shepard, Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming, *Psych. Rev.* **91**(44), 1984, 417–447.
21. R. A. Finke, *Principles of Mental Imagery*, MIT Press, Cambridge, MA, 1989.
22. S. M. Koslyn, *Image and Mind*, Harvard Univ. Press, Cambridge, MA, 1980.
23. P. A. Kolers and H. L. Roediger III, Procedures of mind, *J. Verbal Learning and Verbal Behav.* **23**, 1984, 425–449.
24. I. Biederman, Recognition-by-components: A theory of human image understanding, *Psych. Rev.* **94**, 1987, 115–147.
25. M. J. Tarr and S. Pinker, Mental rotation and orientation-dependence in shape recognition, *Cognitive Psych.* **21**(2), 1989, 233–282.