Perception of Features and Perception of Objects*

Daniel Burnston and Jonathan Cohen†

The binding problem in perception deals with the question of how we achieve the experience of a coherent world of integrated objects, and avoid seeing a world of disembodied or wrongly combined shapes, colours, motions, sizes and distances. In brief, how do we specify what goes with what and where? The problem is not an intuitively obvious one, which is probably a testimony to how well, in general, our brains solve it. We simply are not aware that there is a problem to be solved. Yet findings from neuroscience, computer science and psychology all imply that there is.

— Treisman (1998, 1295)

There is a long and distinguished tradition in philosophy and psychology according to which the mind's fundamental, foundational connection to the world is made by connecting perceptually to features in the environment. On this picture, which we'll call *feature prioritarianism*, minds like ours first make contact with the colors, shapes, and sizes of distal items, and then, only on the basis of the information so obtained, build up representations of the objects that bear these features. The feature priority view maintains, then, that our perception/knowledge of features asymmetrically depends on our perception/knowledge of simple features.

In this paper, we will argue that, the long and distinguished tradition in its favor notwithstanding, feature prioritarianism amounts to an untenable picture of the way our minds are related to the world, and will propose a possible alternative. We contend that there is abundant evidence, drawn from many different aspects of the operation of the perceptual system, that collectively tells against the exclusive priority of feature representations over object representations. Moreover, we believe that this evidence leads in the direction of a more complicated conception of perception according to which our representations of features and representations of objects are highly interactive and mutually informing. Thus, we will argue that perception is most plausibly construed in neither feature prioritarian nor object prioritarian terms, but by an *integrative* view on which neither features nor objects are always representationally prior.

*This paper is entirely collaborative; authors are listed in alphabetical order.

†Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, dburnsto@ucsd.edu, joncohen@aardvark.ucsd.edu

We'll begin, in §1, by reviewing briefly the commitment to feature prioritarianism present in the thought of a number of philosophers and psychologists; then we'll review a range of empirical evidence that, in our view, tells strongly against feature prioritarianism. Next, in §2, we'll turn to the task of providing an alternative general account of perception: we'll set out our favored integrative account, describe its architecture, show how it applies to several areas of current research in perception science, and draw out implications for current debates about modularity. In §3 we'll conclude.

1 Feature Prioritarianism Redux

1.1 Feature Prioritarianism in Philosophy and Psychology

Feature prioritarianism shows up in a variety of roles in a wide range of philosophical and psychological projects; though we can't hope to canvass all of these, in this section we'll run through a few of what we take to be the most prominent reasons for endorsing it.

We begin with one of the most longstanding and widely held motivations for the view within philosophy, which runs through a familiar cluster of empiricist views about the nature of ultimate perceptual justification. Proponents of the views we have in mind have held, first, that perception is the most likely source of foundational justification (i.e., of states that do not require inductive support to be justified) for belief and knowledge, and, second, that the perceptual representations that can best provide justification without depending on further inferential support are those mechanically transduced by causal contact with the world. As it turns out, however, theorists sympathetic to this picture have tended to believe that what minds like ours transduce mechanically from the world are (only) representations of "basic" features — canonically, color, shape, size, and the like. Hence, on this line of thought, feature prioritarianism falls out of the needs of epistemology together with psychological facts about the scope of transduction. Details aside, the point to note for now is just that the kinds of foundationalists we are discussing are committing to the idea that our representations of everything other than the basic features, including complex features and objects, are derived from these initial basic feature representations. Consequently, proponents of this view are committed to the asymmetric dependence of object representations on feature representations: they are feature prioritarians.[1]

[1]That's a little quick. In principle one might deny that the epistemological view under discussion has straightforward implications for psychology, and in particular that the epistemically basic/derived distinction amounts to a distinction between simple and complex mental representations. On second thought, however, this reaction sits uneasily with the appeal to brute transduction as the explanation of epistemic simplicity. It is hard to see why the kind of simplicity a representation gets from being delivered by transduction rather than inference/conceptual combination should count in favor of its epistemic simplicity unless your notion of epistemic complexity follows pretty closely on the heels of your understanding of psychological complexity. We'll put this nicety aside in what follows.

We find (something like) this form of feature prioritarianism present in the thinking of several influential empiricists.

For example, we see some such view in Locke (1975). Recall that Locke distinguishes between simple and complex ideas, and holds that the complex ideas are mereologically constituted from their simple constituents. Moreover, Locke holds that all ideas have their ultimate source in either "sensation" (/perception) or "reflection" (a separate source for simple ideas of, e.g., "perception, thinking, doubting, believing, reasoning, knowing, willing, and all the different actings of our own minds" (*Essay* II.1.4)). But since neither objects nor features seem to be activities of our minds, ideas of either will have to come from sensation rather than reflection. Crucially, however, Locke appears to assume that the only simple ideas that sensation can supply are ideas of sensible qualities — viz., features. It falls out of his view, then, that our ideas of objects are complex, while our ideas of (at least some) features are simple. And since Locke thinks that the acquisition/deployment of complex ideas requires the acquisition/deployment of their simple constituents, he appears to be committed to a fairly strong versions of feature prioritarianism.

A structurally similar form of feature prioritarianism falls out of the logical atomism of Russell (1914). Russell holds that the logical atoms are "little patches of color" and that enduring physical objects are derivative logical constructions from such atoms. Moreover, Russell holds that the epistemology in this respect mirrors the metaphysics: little patches of color are immediately given in sensation, and knowledge of enduring objects is obtained derivatively, via knowledge of the logical constructions from which those enduring objects are assembled out of their atomic constituents.[2]

Indeed, we even find the same basic feature prioritarian picture in the epistemological ruminations of the great anti-foundationalist Quine (1960). For all his anti-foundationalism, Quine clearly thinks there are initial data (which, we suppose, means that other data are non-initial), and that these data just are the irritations of our receptors. But it turns out that what irritates our receptors are not (per se) objects, but features of objects. For Quine thinks that it is only as a response to these irritations/initial data that we go on to posit ordinary objects (and scientific kinds) as a way of coordinating our responses to the

[2]It is tempting to attribute a similarly motivated feature prioritarianism to Carnap (1967), who holds that experience of physical/distal objects is ultimately built from "elementary experiences" (canonically, of colors and other features occurring at visual field points). Likewise for Goodman (1951), who maintains that our experience of concrete individuals is assembled from mereological sums of atomic individuals (color qualia, times, and visual field places). However, this reading of Carnap and Goodman is inconsistent with their repeated, explicit insistences that the dependency relations they unpack by the abstractive, reconstructive process of "quasi-analysis" is without epistemic, psychological, or representational import (Carnap (1967, 109, 149), Goodman (1951, 151)). While, in our view, such insistences make it much less clear just how to understand the nature of the "constructional" priority relations these authors take themselves to be elucidating, it also means that it would be inaccurate to assimilate their views to the kind of epistemic motivation for feature prioritarianism under discussion. (Thanks to Robert Briscoe for correcting us on this point.)

impinging array of receptor irritations. Yet again, the feature prioritarianism could not be clearer: it is features first, and objects second.[3]

Although we could on with further examples, we propose to take it as read in the following that feature prioritarianism is a position that has been attractive to many, and that goes to the heart of some widely known and justly influential conceptions of epistemology and perception.[4] However, feature prioritarianism also plays important roles in other, less epistemologically oriented philosophical debates. One of these is a range of discussions concerning the "units" of perception. Siegel (2006) discusses a variety of views of what gets represented in perception, categorizing as "minimalist" views that would only posit representations of such things as colors, shapes, and illumination, without performing any binding operation into objects, etc. More "permissive" are views that eventually categorize such feature representations into object representations or, yet more permissively, representations of kinds and categories (her preferred view). While not explicitly committed to the idea that object and kind representations are composed from feature representations via a mereological or logical operation, a form of feature prioritarianism is present here: it is the idea that any view of perception must posit representations of features, and only more permissive views allow that we represent such things as objects and kinds at all. This amounts to a methodological variety of feature prioritarianism, in that it carries an implicit commitment to the view that representations of sensory features comprise the most basic way of perceptually accessing the world.

There are, to be sure, figures who reject this methodological form of feature prioritarianism. Thus, Orlandi (2010), drawing on "embedded" views of perception, argues that perception skips the representation of features entirely, and gains access to the world via a direct relation to objects. And Schiller (2012) argues from cases of perceptual illusion and attention effects that the primary units of perception are actually facts or states of affairs, in which

[3]As the discussion of Quine should make clear, feature prioritarian presuppositions are not confined to those with classical foundationalist sympathies, even if such presuppositions are often more explicit and easier to find in foundationalist writings.

[4]We emphasize that, despite their prominence and historical importance, the epistemological motivations described above comprise but one line of support for feature prioritarianism — and, crucially, one that comes with idiosyncratic features from which we'll want to abstract below. Specifically, those theorists whose commitment to feature prioritarianism comes only or mainly from the role they assign perceptual states in foundationalist forms of justification are naturally understood as interested in only those perceptual states that are consciously accessible — after all, consciously inaccessible states are by definition unavailable for citation in subjects' justifications. (Or, if you like, these theorists can be construed as concerned with all perceptual states, but only on the assumption that perceptual states are uniformly accessible to consciousness.)

In contrast, our main concern in this paper is with the broader question of the relation between feature perception and object perception in perceptual systems quite generally, without limitation to that subset of perceptual states that are consciously accessible. Thus, despite our claim that the view is motivated in part by epistemological motivations, we won't be construing feature prioritarianism (or the alternatives to it we'll explore later) as limited to conscious perceptual states, and will feel free to draw on unconscious states in the critical discussion that follows. (Thanks to Chris Hill and Robert Briscoe for prompting us to clarify our stance on this point.)

both properties and objects are represented relative to a context. However, the very extent to which these thinkers are (self-consciously) aiming to overthrow received opinion in favor of a radical alternative view of perception shows how deeply feature prioritarianism is intertwined with contemporary thinking about perception.

In addition to these philosophical debates, feature prioritarianism plays a significant role in a variety of psychological and neuroscientific investigations, which standardly assume a roughly feature prioritarian organization in perceptual systems. In such discussions, perceptual features are typically described as occurring at lower levels of a representational hierarchy, while objects are represented in higher levels, where, vitally, the object representations result from computations performed over the feature representations. Generally, the simplest features, such as light discontinuities (edges) and luminance, are posited to be represented at the lowest levels, while intermediate features such as shape, color, and motion are represented at intermediate levels, and form the direct basis for object representations. Thus, perceptual hierarchies are committed to a relatively strong form of feature prioritarianism.

One field in which the perceptual hierarchy idea has proved important is visual neuroscience, in which the feature hierarchy view has provided one of the predominant conceptual guides to empirical investigation of the visual system over the last 30 years. Physiological evidence from a number of methodologies has provided evidence that the visual system is indeed sensitive to more complex features at later levels of visual processing, but theorists differ in how explicit and independent the representations at different levels of the hierarchy are posited to be. A particularly strong view is posited by Zeki (2001); Bartels and Zeki (1998) who argues that not only are feature representations entirely explicit at each particular level of the visual hierarchy, but the representations at each level are sufficient to produce conscious visual experience of these features. Moreover, there is a philosophical motivation for this sort of view, as Zeki argues that gaining knowledge of perceptual features is a primary function of perceptual systems. While other views are less strongly phrased, a glance through the classical perceptual neuroscience literature shows that feature prioritarianism is strongly present in the views of a large number of investigators.

Finally, feature prioritarianism is an inherent part of standard views of object recognition and classification in computational psychology, where it is almost universally assumed that object recognition occurs by combining feature representations (Westphal and Würtz, 2009). Theorists differ on whether the feature representations are general (usually shape features; Biederman (1987)), or whether they are specific to the object-class being recognized (Ullman, 2007). Particular models also differ in whether they are purely bottom-up (Riesenhuber and Poggio, 1999) or employ top-down feedback (Epshtein and Ullman, 2005), but in either case it is particular, intermediate level features (and their arrangement in the scene) that are used to recognize objects. Thus, the perceptual hierarchy view, and its concurrent feature prioritarian leanings, is an important conceptual and explanatory construct for both neuroscientific and psychological investigation.

The considerations above suggest that a commitment to feature prioritarianism is shared by a large number of theorists with otherwise disparate views, and that this commitment plays a fundamental role in their thinking about perception. This makes all the more important the question of whether that commitment is true. In what follows we'll suggest that it is not.

1.2 When Features Fail: Object Representations Without Feature Representations

In this and the next few sections we'll review a number of findings that, in our view, tell against the feature prioritarian insistence that perception of objects is always asymmetrically grounded in prior perception of features. We'll try to make our case against feature prioritarianism in two ways. First, in §§1.2.1–1.2.2 we'll consider a set of techniques researchers have developed for studying object perception and argue that, taken collectively, they suggest that the perception of features is not always necessary or semantically sufficient for the perceptual representation of objects. Then, in §§1.3.1–1.3.2 we'll turn to work on the representation of features; here we'll be suggesting that, at least in some cases, and contrary to the feature prioritarian's proposed dependency ordering among representations, it looks as if our representations of features are mediated by prior representations of objects.

Before we begin to review the empirical evidence, however, five caveats about the terms of the discussion are in order.

Our first caveat concerns the notion of objects we have in mind. Famously, the term 'object' in cognitive science is ambiguous between a "thick" reading on which it applies to ordinary physical objects (cabbages and kings, rocks, trees, and flowers) and a "thin," more specialized usage that applies to (roughly) bounded and connected things that trace out continuous paths in spacetime. One crucial difference between the two notions comes out in their persistence conditions: thick objects, but not thin objects, must retain their kind affiliations to count as persisting over time. In what follows, when we talk about the perception of objects, we mean to be talking about objects in the thinner sense. Our discussion below requires only a relatively minimal understanding of thin objects, and we prefer to remain as uncommitted to further details as we can.[5] Minimally, we need it to be true of objects that they can be distal, that they can bear features that perceptual systems can be sensitive to, that they can serve as the loci for feature binding, and that they can be served up by perception qua the referents of demonstrative thought and language on at least some occasions. Beyond that, we invite readers to rely on their own preferred understanding of what thin objects amount to.

[5]Spelke (1990) develops what has become a prominent way of thinking about the nature of these thin objects towards which we are broadly sympathetic. Despite this sympathy, however, we don't intend our view to be committed to the particular details of Spelke's account.

A second point is that, as we are understanding it, feature prioritarianism is a position about representational rather than metaphysical priority. Accordingly, the issue on which the position hangs is not whether features or objects are primary qua elements of the mind-independent world. It is, rather, whether features or objects are *representationally* basic (for beings like us). Specifically, we take feature prioritarians to be committed to the view that all of the semantic resources needed for (justifiably) representing objects are present in representations of features. It follows from this basic commitment not only that feature representations must be temporally and causally prior to object representations, but that the latter can be formed by arranging the materials made available in feature representations in the right way — where what counts as the right way can be a logical or mereological relationship.[6]

A third caveat concerns the understanding of representational priority itself. It should go without saying that if subjects entertain object representations, their doing so is caused by some event or state of the world, broadly speaking, rather than by magic. Therefore it is trivial that there will always be an instantiation of some or other cluster of features that causally contributes to whatever token object representations in fact occur. On the other hand, what this doesn't show is that the subject is tokening her object representation as a result of first representing the features in question — she may never represent them, or may represent them later, or optionally, or whatever. A fortiori, the trivial point that there is a featural *causal* basis for every object representation can't be taken to show that every object representation is a logical or mereological complex with representations of the relevant features as its constituents. It is this last issue — and not the trivial question of whether there is a featural causal basis for object representation — that we mean to be joining in the following.

Here is a fourth caveat about the dialectical situation. We have intentionally constructed as our target a very strong version of feature prioritarianism — one offered as a universal generalization about the psychology and epistemology of perception. And we will argue against that view (in the rest of this section) by focussing on cases of perceptual phenomena about which that universal generalization is false. But this argument leaves open the possibility that there are some perceptual phenomena about which feature prioritarianism is true. (Nonetheless, we will suggest that the types cases we cite are sufficiently robust to motivate a search for a better general view of perceptual architecture.) Moreover, in urging rejection of the strong form of feature prioritarianism, we do not mean to insist on abandoning all the projects with feature prioritarian presuppositions. Plausibly, some explanations relying on feature prioritarianism might survive its rejection *qua* overall view of

[6] We mention this to prepare the ground for our argument against the view in §§1.2.1–1.2.2, which will consist in finding cases where object representations outstrip the resources made available by feature representation. Effectively, then, we'll be arguing that (in the cases we mention) perception is ampliative, and that what it adds cannot be accounted for by merely logical and mereological constructions from feature representations.

perception, though sorting this out would require case by case assessment in each domain.

Our final caveat is that we do not suppose ourselves to be the only writers who have recognized the import of the cases we cite; nor do we take these particular cases to exhaust the class of relevant examples. Indeed, several perceptual scientists have seen the relevance of the cases we discuss for thinking about the computations performed by perception, and have come up with insightful algorithmic ways of thinking about them (Adelson, 2000; Yuille and Kersten, 2006). Moreover, there are episodes in the history of psychology during which many theorists have held anti-feature-prioritarian views.[7] That said, in our view, the historical pendulum has generally swung in a more feature prioritarian direction since the rise of computational approaches to vision and perceptual neuroscience. Overall, then, we find feature prioritarianism to be a currently popular framework for understanding perception, despite the known existence of cases that, when taken seriously, tell against it. We surmise that this persistence is due both to feature prioritarianism's intuitive appeal and its long history. In what follows, then, our modest aim is to clarify the nature of the challenge posed to feature prioritarianism by these cases, and to provide a schematic alternative view of overall perceptual architecture that avoids these problems while, perhaps, beginning to reconcile some of the historical fissures we've mentioned.

1.2.1 Objects Without Features: Apparent Motion

We begin by considering a first technique for studying object perception that turns on apparent motion phenomena. In one version of this type of phenomenon, the experimenter presents two images in succession: the first contains a line on the left hand side and nothing on the right, and the second contains a line on the right hand side and nothing on the left. At certain rates of succession (and over certain spatial distances between the lines in the two images), subjects perceive (not two static lines, but) a single line that moves from left to right.[8]

[7] Thus, Robert Briscoe has suggested to us that several figures within the Gestalt tradition (e.g., Koffka, 1935) may be examples.

[8] A range of different effects are standardly grouped under the "apparent motion" heading. In a historical discussion of the original work on apparent motion in the early 20th century, Steinman *et al.* (2000), distinguish between the Φ-effect, which occurs at short frequency temporal alternations, and does not involve perception of an object in motion, and the β-effect, which occurs at longer frequency temporal shift and does involve object perception. However, Braddick (1974) discusses short temporal and spatial displacements in dot-motion perception, and seems to suggest that even at the short temporal range, objects can be tracked via apparent motion. We will abstract away from these considerations, as the distinctions here are less than clear in the current literature. When we discuss "the apparent motion phenomenon," we refer only to the longer frequency temporal alternations, which, as we will show, involve object perception. (Of course, this should not be taken to rule out other sorts of cases from counting as apparent motion.)

Matthen (2005, 278ff) also discusses apparent motion effects in arguing against exclusively feature-based views about perception.

Apparent motion phenomena provide clues to the visual perception of objects because in such cases the (illusory) perception of motion happens when perceivers perceptually (erroneously) represent two spatiotemporally separated regions as falling in the spacetime trajectory of a single object. Therefore, by studying the range of conditions under which illusory motion can be induced, we gain insight into an important class of constraints on object perception. In particular, these cases are fertile grounds on which to test feature prioritarian requirements on object perception.

What is particularly interesting about apparent motion phenomena for our purposes is that, perhaps surprisingly, they do not require persistence of shape, color, size, and other features (Kolers and Grunau, 1976). On the contrary, it is possible to induce apparent motion between objects that differ in these features. For example, if we replace the line in the second image of the classic Φ demonstration with a circle, then successive presentation (within the necessary time interval, over the necessary distance) of the two images — i.e., first a line on the left hand side of the visual field, then a circle on the right hand side of the visual field — results in perception of a single object that changes its shape as it appears to traverse the space between the left and right sides. Likewise, successive presentation of objects that differ in color or size result in the perception of a single object that changes its color or size. Taking these results at face value suggests that the perceptual representation of objecthood does not require the persistence of a particular color, size, or shape.

The vital feature of this type of phenomenon for our argument is that it speaks against the representational sufficiency of feature representations. It does so in two ways. First, since the phenomenon can occur over a large range of features, which stand in different relationships to each other in the different trials, it suggests strongly that there is not a particular set of features whose representation is necessary for generating the object representation. Second, and perhaps more compellingly, the spatial and temporal discontinuities between the changing features show that neither the logical nor mereological construction views are sufficient to generate the object representation. On the side of logic, it is certainly logically compatible with the presented stimulus that there are in fact two objects being presented (and, moreover, this is true!). Similarly, the represented object is clearly not mereologically composed from perceived features — after all, the object representation persists over discontinuities in those features. Thus, the resources of the putative feature representations are insufficient — there is a gap between them and the perceived phenomenon that no logical or mereological construction can fill.

To repeat an earlier warning, we are not contending that the perceived features play no role in the process of object perception here. It is clear that the features presented stand in some relations to the eventual perception. For one, the presented features play a (partial) causal role in bringing the perception about. There also seems to be a good sense in which the presented features are an informational resource for perception — they carry information about the distal world that is exploited by perception. We will say much more about this in the ensuing discussion. What is important for now is that neither of these relations is sufficient to generate the perceptual phenomenon: more must be going on. Thus, even if there may be features that have a causal or information priority with respect to the eventual object representation, this falls far short of the feature prioritarian claim that there is a representational priority between the two.

Now, the findings about apparent motion do not, by themselves, refute feature prioritarianism. For one thing, feature prioritarians might insist that object perception is here mediated by the perception of features *other* than color, size, or shape; in effect, this would be to posit *further* features involved in the scene that fill the gap between the clearly presented features and the eventual perception. We cannot rule out this response: the results we've cited do not and cannot show that apparent motion can be induced between presentations that differ in respect of any arbitrary feature. They do, however, show that apparent motion (hence objecthood) can be induced between objects that don't share the very features that actual feature prioritarians have explicitly put forward as underlying object perception, and so tell against canonical versions of the view. The explanatory burden now shifts heavily to the side of the feature prioritarian, who must delineate a feature other than the size and shape of the presented lines, explain why it is clearly represented in the scene, and show that its representation accounts for the representation of objects in the scene. We confess that we cannot think of any such features. The only other feature perceived is motion, but the perception of this feature seems to be intimately tied up with the perception of the object — i.e., one switches from seeing successively presented lines or shapes to seeing an object that moves. Thus, we view the present response as a slim reed for the feature prioritarian.[9]

Here is a more interesting line of response. A feature prioritarian might try to solve the problem not by positing a specific further feature that is present in the scene and compels the object interpretation, but instead by insisting that a certain arrangement of features is all there is to the object representation. Alternatively, but in a similar vein, she might focus on the diachronic nature of apparent motion, and import further facts about feature temporality and location in order to account for the object representation solely on the basis of represented features.[10] One of these related moves will always be available to the defender of feature prioritarianism (which one will depend on the nature of the case), and there is nothing in the apparent motion

[9] Needless to say, feature prioritarians are welcome to pin their hopes on other features; whether these will succeed or not is not something to be settled from the armchair. However, in line with our general pessimism on this score, we note that there is reason to think that no visual features will account for all cases of object representation — even for visually presented items. Thus, for example, Shimojo *et al.* (2001) show convincingly that whether two visually presented shapes that cross in spacetime are interpreted as colliding or bouncing off one another (viz., how the visual system parses the scene into objects) depends on whether there is an auditory ping presented just at the moment of intersection. Again, the feature prioritarian can accommodate these results by suitably modifying her inventory of the features on which object perception depends. Whether there is any fully adequate version of the position (and whether it will bear any resemblance to the feature prioritarianism with which its proponents began) remains to be seen.

[10] Thanks to Christopher Hill and Christian Wüthrich for pressing the objection based on the diachronic nature of the case to our attention.

phenomena (or other cases we will present) that logically rules them out. However, the opponent of feature prioritarianism has an equally generalizable response. For the current "arrangement of features" suggestion comes with no principled account of which features get bound into objects in which cases — that is to say, on the current view there are no explanatory resources for privileging the particular arrangements of features, or particular spatial and temporal locations of features, that get bound into objects. As we've seen, apparent motion occurs over a very broad range of features, across gaps of spatial continuity between those features, etc. The defender of the sort of objection under consideration therefore needs a different set of features and arrangements thereof to account for each particular case; but this way of proceeding is objectionably unsystematic and ad hoc. As such, we view this general type of defense of feature prioritarianism as a non-starter.[11]

1.2.2 Objects Without Features: Object Tracking Through Occlusion

A second technique for studying object perception involves visually tracking objects as they move about a computer screen.[12] It turns out that subjects can successfully track 4–5 such objects as they move randomly and independently among a larger set of identical, also randomly and independently moving items. Indeed, subjects can continue to track these items even if they change their features (e.g., color, shape, size) while in motion. For reasons raised in connection with our discussion of apparent motion, we believe the persistence of object tracking through changes in such features is *prima facie* evidence against feature prioritarianism: here, too, subjects take there to be one object that changes its features, which is not what one would expect if one thought they were arriving at their object representations exclusively by inferences from evidence about the features exemplified before them.[13]

[11]A corollary of this point applies to those tempted to defend feature prioritarianism on epistemic, as opposed to merely psychological, grounds. As we saw in §1.1, some feature prioritarians have held that perception of features is not just temporally prior to the perception of objects, but that representations of features are formed first and then *serve as evidence that warrants the attribution of objects*. These feature prioritarians believe that perceptual representations of features are not only temporally prior, but also psychologically and epistemically prior to representations of objects. But if this is the correct picture, then it seems to us that the cases of apparent motion under consideration are genuinely puzzling. For, in such cases, subjects do seem to be picking up on *two* color (/shape/size) properties rather than one — this is why they describe the moving objects as morphing in color/shape/size along the trajectory of apparent motion. But if they pick up on two such features rather than one, and if features are indeed serving as the data/premises justifying an inference to the presence of objects, one would predict that subjects would (correctly, as it turns out) infer the presence of two static objects rather than (incorrectly) inferring the presence of one moving object. That perception does not accord with this prediction in at least some cases, therefore, tells against this version of feature prioritarianism.

[12]For an overview, see Scholl *et al.* (2001); Pylyshyn (2003, 2004).

[13]A further, and particularly elegant, study involving tracking of multiple objects through featural change is described by Blaser *et al.* (2000). These authors presented subjects with two superimposed (spatially coincident) Gabor patches, each of which had at each moment a color, a spatial frequency, and an angular velocity, and showed that subjects could track one of the patches (as it might be, one that is initially white, narrow in spatial frequency, and rotating clockwise),

However, there is an additional finding from this research that we find even more striking in the present context: it turns out that subjects can continue to track objects through occlusion "behind" other objects on the screen (Scholl and Pylyshyn, 1999). On its face, this finding is hard to square with the feature prioritarian insistence that object perception always depends asymmetrically on feature perception, since in the cases at issue perceivers appear to be representing objects even when the features of those objects are unavailable to perception (because they are entirely occluded from view).

We suspect that a dedicated feature prioritarian would attempt to save her view from this concern by suggesting that the object representations in such cases are, after all, derived from the perceptual representation of features, but features *other* than those of the occluded objects. Thus, for example, a feature prioritarian might suggest that the object representations are, in such cases, derived from the features of the occluders. Unfortunately, this solution won't cover all the cases; for it turns out that subjects continue to track objects through occlusion even if the occluders are themselves invisible — i.e., if the occluders are entirely unmarked regions where intersecting objects become invisible as they pass through and then become visible once again after emerging (Scholl and Pylyshyn, 1999). Since, in the latter sort of case, features of the "virtual" occluders are similarly unavailable to perception, it's hard to see how these features could serve to ground the object representations we see.

We expect that the feature prioritarian will reply, at this point, that the subjects' representations of persisting objects during occlusion are grounded in features that were perceived earlier than the occlusion. The idea would be that the perceiver represents a cluster of occurrent features (a shape, a size, a color) at time t_1 earlier than the occlusion, and that her subsequent representation of an occluded object at later time t_2 is indeed dependent on (computed/inferred from) features, as per feature prioritarianism — only that the relevant features in question would have been perceived at t_1 rather than t_2.

We are in fact extremely sympathetic to something like this suggestion as a general (if loose) characterization of what goes on in multiple object tracking. However, it is not clear that the idea can save feature prioritarianism. For one thing, some have worried about whether our cognitive architecture can bear the computational load of maintaining and updating information about the identity of tracked objects on the basis of features alone (Pylyshyn and Storm, 1988).

A more pressing worry for the present suggestion parallels the considerations raised in the last subsection — namely, the question of what legitimates the inference from the earlier features to the later object. To see the force of the worry, note that it is not generally true that representation of features exemplified at an earlier time leads us to represent an object at a later time:

while both patches simultaneously varied independently and randomly in these features. Like the other results we are reporting, this experiment shows that subjects can indeed track items (viz., maintain object representations) through featural change. But this experiment also shows that individuating the objects *need not depend on locational information* — that wouldn't distinguish the items in this case, since they are spatially coincident.

perceiving the characteristic shape and color of the banana at this morning's breakfast table (by itself) does not warrant a representation of a banana on the lunch table at noon. Moreover, as we just saw, subjects do not assume that objects always retain their features, since objects are classified as such across changes in those features. In contrast, the point about tracking through occlusion is that the later representation of an object *is* computed from (/justified by) earlier representational states; it is just unlikely that feature representations play this role.

What, then, would computationally/justificationally legitimate an inference from features at t_1 to an object at t_2? In a word, objects. That is, the inference to a later object representation looks much better if it is made on the strength of earlier representations of both features and objects (or, better, features bound to objects). This is because perceivers know that objects tend to follow continuous trajectories in space-time (Spelke, 1990), and so can and do infer from earlier object representations to later object representations. It would seem, then, that a feature prioritarian can, after all, hope to explain away the initially troubling cases of tracking through invisible occlusion by appealing to representations formed prior to the occlusion. It is just that the cost of doing so is accepting, contrary to feature prioritarianism, that the needed earlier representations include representations of objects that are not simply combinations of features.

1.3 When Features Fail: Object-Dependent Feature Representations

Until now we have been reviewing evidence intended to show that (contrary to feature prioritarianism) there can be perceptual representation of objects without perceptual representation of features. We now want to argue that in some cases the direction of dependence runs opposite to that urged on us by feature prioritarians — that in these cases the perceptual representation of features is mediated by the perceptual representation of objects.

1.3.1 Object-Dependent Feature Representations: Perceptual Constancy

One canonical sort of case that arguably has this structure involves perceptual constancy. Thus, consider the following kind of perceptual scenario, depicted in figure 1. You are viewing a coffee cup under uneven illumination; suppose there is an illumination edge falling across the cup, so that half of it is in direct sunlight and half of it is in shadow but is illuminated diffusely. In this scenario, much more light is reflected to the eyes from the directly illuminated than the diffusely illuminated portions of the coffee cup. Now, we know (e.g., from matching/sorting data) that the there is some good respect in which subjects represent the adjacent but differently illuminated regions of the cup's surface as sharing a color (a feature).[14]

This is, of course, a perfectly textbook instance of perceptual constancy, wherein the perceptual system assigns distal features (in this case, color) in a way that is unvarying even though there is variation in the perceptual signal. But we now wish to point out that, at least in some cases, those feature assignments are mediated by information regarding the objects there are in the scene. In fact, this is not surprising with respect to the particular case under discussion: it is to be expected that we more easily treat the difference between the two adjacent cup regions depicted in figure 1 as deriving from an abrupt luminance change (so not a change that will matter to our assignment of stable surface properties to those regions) if we start out with the assumption that the regions belong to one materially homogeneous object. The point, then, is that what features are assigned to the regions of the perceptual scene depends on just how the scene is parsed into objects. And this is just to say that (exactly contrary to what feature prioritarians claim to be universally true), it is a case in which feature representations depend on object representations rather than the reverse.[15]

A feature prioritarian might object at this point that, contrary to what we say, her view can in fact accommodate such cases of perceptual constancy if supplemented by reasonably standard and independently well-motivated assumptions about the architecture of perceptual processing. On the imagined response, the feature prioritarian would begin by holding that features are separately and independently represented, and lead to object classification via being combined by a further processing stage. However, she could go on to claim that in many cases the object itself, here the coffee cup, is recognized by some subset of the potential features available in the scene — for instance its shape — and this recognition then in turn provides feedback to the feature-recognition level, modifying the feature percept accordingly. If so, would this not show how feature prioritarianism can accommodate the sorts of cases we are bringing to bear?[16]

We have two concerns about this response. First, notice that the shape of the relevant object in a scene like that depicted in figure 1 is not simply apparent in the figure — i.e., it is not simply a number of easily perceptible edges in the shape of a cup. The edges of the cup themselves consist of a variety of different

[14]This characterization is somewhat more cagey than the usual claim that, under this scenario, perception represents the two regions as alike in color. We favor the less simple characterization because there's also plenty of empirical support for the claim that perception represents the two regions as (in some other good respect) *unlike* in color, and because the simpler characterization can only stand if we have reasons to reject the second sort of representation as being irrelevant to color (for discussion, see Cohen, 2008, 2012). We leave these issues aside in the main text.

[15]We emphasize that object-mediated perceptual constancy shows up for exactly the sorts of features — color, shape, size, etc. — that feature prioritarians have traditionally held to comprise their privileged circle of basic features. Therefore, the phenomenon is damaging not only to the universal dependency claims made by feature prioritarians, but also to the specific instances of those dependency claims that involve feature prioritarians' historically preferred featural basis.

[16]Thanks to Wayne Wu for pressing us on this point.

wavelength and luminance contours, which stand in a range of relationships to each other. This suggests that there is not a clear processing pathway that goes from shape to object, and then feeds back to color, because the features that are supposed to underlie the object classification on the current view are themselves (partially) dependent on the very features that are modified. Far from a clear bottom-up and top-down arrangement, it seems that the features themselves, along with the object interpretation, interact in processing the stimulus. (This is just the sort of view we'll be defending in §2.)

Our second concern about the proposed alternative is computational. The imagined feature prioritarian's proposed top-down explanation of constancy effects posits that there is something like an object-recognition module that takes in representations of features, produces the appropriate object representation, and in turn modifies the original feature representations according to that representation. Crucially, this proposal relies on positing an extremely powerful module for object-recognition — one that can keep track of all of the relevant feature representations, and knows how to modify each feature representation in response to the object recognitions it produces. And while (once again) nothing logically rules this view out, we consider it unlikely that the wide range of perceptual effects under discussion here are the result of a single module or object-processing level.

1.3.2 Object-Dependent Feature Representations: The Cornsweet Illusion

We invoked perceptual constancy to show that there are cases where the perceptual representation of features systematically depends on the perceptual representation of objects. Specifically, the cases we raised were those in which, because of its object assignments, the perceptual system represents featural sameness across variation in stimulus values. But now we want to point out that object assignments can have the opposite effect of causing the perceptual system to represent featural difference across sameness in stimulus values.

A standard case that takes this form is the so-called Cornsweet illusion (also called the Craik-O'Brien-Cornsweet illusion), depicted in figure 2 (O'Brien, 1959; Craik, 1966; Cornsweet, 1970). In this configuration there are two narrow opposed lightness gradients separated by an edge; on the far sides of the lightness gradients are two extended patches of equal lightness. It is a robust finding that subjects perceive the two equal lightness patches toward the outsides of the stimulus as differing in lightness, with the one near the higher lightness area being perceived as lighter than the one near the lower lightness area.

What is interesting for our purposes is that the way in which lightness features are assigned in this configuration appears to depend crucially on the way the scene is perceptually parsed into objects (Purves *et al.*, 1999). Thus, for example, the effect is enhanced when there are sufficient depth cues for subjects to interpret the scene as involving a surface, and the edge as a curve or break in that surface. Conversely, if the patches are displayed as vertically orientated, with the light-to-dark gradient on top, the effect is decreased

Figure 1: There is some good sense in which the regions of the cup in shadow and the regions of the cup in direct sunlight look the same in color. Photograph © Jonathan Cohen.

Figure 2: The Cornsweet illusion.

considerably, as this presentation is more consistent with the interpretation that the two patches are on the same face of a surface (with light coming from above). Also importantly, the effect is virtually eliminated if the contrasting background normally presented with the figure is replaced with a uniform background, thus eliminating the surface interpretation. As in the case of perceptual constancy, then, these cases suggest strongly that the perceptual representation of features (here, lightness) can depend systematically on the perceptual representation of objects, *pace* feature prioritarian insistences to the contrary.

Now, as before, we must admit that the feature prioritarian could respond to this challenge by coming up with further features, present in the scene in question, that might compel (or legitimate an inference to) the eventual perception. Specifically, the feature prioritarian could insist that it is not objects or surfaces per se, but instead a complex arrangement of lines and luminances that causes the Cornsweet illusion. But we think this is something of a desperate move. We have presented a number of diverse situations in which object perception seems either to outstrip or modify feature representation. While the feature prioritarian may be able to handle one or another of these cases individually by pointing to features, it seems to us highly unlikely that she will be able to come up with a unified set of features or arrangements capable of accounting for the perceptual effects in all of diverse range of cases we have surveyed. Instead, the sets of features and arrangements would have to be constructed specifically to cover each individual case. This sort of response seems not only *ad hoc* and unsystematic, but also far removed from the original feature prioritarian vision, which hoped to ground object perception in a few basic features and a certain type (logical or mereological) of arrangement. Once the move to a diverse set of features has been made, a lot of territory has already been ceded to the opponent. Further, as before, such an opponent is now in a position to ask: why do these particular sets of features become relevant in these particular cases? What is the unifying principle in these cases that brings these features together in just these ways? As we saw in the tracking-behind-occlusion cases, the obvious answer is that these features are the ones that aggregate around *objects*. But, of course, this simple and obvious response is not available to the feature prioritarian, and we are at a loss to imagine an equally general and systematic alternative.

1.4 General Lessons

While we have abstracted away from important details, we take the foregoing examples to suggest strongly that the perception of features does not, as a general matter, depend asymmetrically on the perception of objects. The apparent motion and object tracking results show that, in many cases, the perception of objects outstrips the representational resources present in the features perceived. Moreover, perceptual constancy cases and the Cornsweet Illusion demonstrate that the perceptual representation of features is sometimes itself dependent on the perceptual representation of objects. In short, we claim, these findings make it very hard to see how feature prioritarianism could amount to a general and descriptively adequate conception of the mind's perceptual relation to the world.

Why, then, has it proven so popular? After all, the cases we've appealed to in raising problems for the view are anything but unknown to perceptual psychology. We suggest that the popularity of feature prioritarianism rests in part on its having become entrenched through its historical connections with a very large range of empirical projects in the study of perception (§ 1.1), and in part in its being a deeply intuitive view about perception. In order to leave feature prioritarianism behind then, we need an alternative general architecture for perception. In § 2, we sketch the preliminary outlines of such a view, which both accounts for the cases we've discussed and, we contend, aligns better than does feature prioritarianism with a variety of productive projects in perceptual science.

If our arguments in the current section have been successful, then such a view must avoid a commitment to feature prioritarianism. On the other hand, neither do we want to endorse an opposite, "object prioritarian" view, on which the perceptual assignments of features would always asymmetrically depend on the perceptual assignments of features. For while much in this area remains controversial, it is hard to see how subjects could identify all objects without relying on some feature instances — particularly at an early stage of perceptual processing — such as figure/ground segmentation, intensity discontinuities, edges, spatial frequencies, and the like.[17] The sort of view we're after, then, should capture the idea that there is no universal, unidirectional dependency relation between features and objects — it must, in this sense, be a *no-priority* view.

In the remainder of the paper, we offer one such view (others are, of course, possible). In place of the versions of prioritarianism, our proposed alternative will focus on the *integration* of perceptual information. Instead of a representational dependency between features and objects (or vice versa), we submit that a fundamental principle of perceptual function is that perception (and elements thereof) integrates the different types of information available to it (them), in order to arrive at the most coherent interpretation of the arrangement of features and objects in the world. Thus, in general, neither features nor objects are represented prior to each other, but both are represented together as the best interpretation of a given set of perceptual information — that is, features and objects are perceived as a package deal.

[17]There is, of course, much more to say in assessing object prioritarianism. We content ourselves with these extremely brief remarks mainly because the view has been far less popular than feature prioritarianism. In general, even those sympathetic to object prioritarianism recognize the informational/causal importance of features that we have stressed in scene perception. For a more extended argument supporting the representational importance of features — i.e., that perceptual systems need to represent features as well as objects—see Clark (2004), Cohen (2004, §3.3).

2 Perception as a Package Deal: An Integrative View

Here is is the crude picture we have in mind. Imagine that, in the beginning, the world is a blooming, buzzing confusion, populated with all manner of objects and features. The perceptual system finds itself in this world, and is confronted with an initially unstructured or minimally structured (transduced) input signal. The task for the perceptual system, then, is to impose structure on the impinging signal in order recover information about the distal world. The system uses a variety of techniques, or "perceptual strategies" for carrying out this task, where each strategy operates on a particular set of (perhaps various types of) information, and outputs a working model of the arrangement of the particular features and objects that fall within its domain. As the process continues, perception begins to integrate the different working models, for example using its working model of object locations and/or motions to guide its model of the colors and shapes in the scene, and vice versa. Thus, the perception of features and of objects generally occurs in tandem. Feature and object representations are not the result of dedicated processes that deal solely with either features or objects, but instead emerge as a package deal from a variety of different strategies for integrating different kinds of perceptual information, and the further integration of these different strategies into a unified percept.

Now, reconsider the examples we discussed in §1. An integrative view based on perceptual strategies allows for a unified approach to accounting for these cases, since it posits that perceptual information of different types can be relevant for producing perceptual phenomena, and that different types of information can be integrated in different ways, depending on the perceptual strategy being employed. Thus, phenomena of apparent motion are interesting because of the light they shed on an important subset of the strategies subserving perception: the range of conditions in which apparent motion can be induced are windows into the types of information to which those strategies are sensitive and the types of integration they perform. In this case, the relevant strategy seems to privilege information regarding spatio-temporal patterns of stimulation. Though these patterns may be instantiated by features in each case, the strategy does not seem to depend on specific, persisting types of features. After its operation, this strategy might then be integrated with further strategies in order to produce a percept of a moving object with such-and-such features (and, as discussed, these features may shift during the course of the stimulus; on the integrative view this would be the outcome of a cross-temporal series of integrations based on new information). Similarly, on the integrative view, perceptual constancy and the Cornsweet illusion are indicative of perceptual strategies for making sense of the layout of luminance patterns in the visual field by incorporating information about probable object/surface layouts. Focusing on integration thus allows a theory of perception the flexibility to account for the range of perceptual phenomena that feature prioritarianism is at a loss to explain.

2.1 Perceptual Strategies and Inputs

We said that, on our view, perception proceeds by integrating the outputs of a variety of perceptual strategies. Each of these can be thought of as a computation (or series of computations) performed by a perceptual system, and which responds to a specific set of inputs, by producing a stable percept of a (putative) entity or arrangement of entities in the world. (Here we use 'entity' broadly to talk about a feature, object, or set of objects.) What counts as an input for such a perceptual strategy is any source of information about the perceptual scene that is relevant for determining the eventual entities perceived.

One of the things that distinguishes the integrative view from object and feature prioritarianisms is the claim that inputs can come from a variety of sources, and will vary depending on the strategy being employed. Each is potentially equally important (within the strategy) in determining the nature of the computation and the eventual percept produced. As such, and as its name is intended to suggest, the integrative view is not committed to any robust unidirectional representational dependency between features and objects.

The integrative view also easily accommodates two further (well-supported) ideas about inputs. The first is that inputs needn't be solely exogenously generated. Rather, the view welcomes the idea, which has become a standard tool in current perceptual science, that perceptual systems store information or assumptions about the types of environments in which they are employed. On the integrative view, such stored information also serves as inputs to a given perceptual strategy.[18] Second, there is no reason, on this view, that inputs to a computational process should be restricted to a single sensory modality (we have already observed one instance of a perceptual strategy that appears to work on information from distinct modalities — see note 9).

Together, these points suggest that the class of perceptual strategies is both multifarious in its computations and graded in complexity. It is multifarious in that different perceptual strategies use a variety of different types of information, in different combinations. Perceptual strategies can also be more or less complex depending on the number of different types of information

[18]Thus, to use a canonical example from Ullman (1979), the perceptual system has enough information to compute three-dimensional structure from motion only if it works on the assumption that bodies in our environment are typically rigid. In §2.2 we mention further examples of the appeal to such assumptions involving color perception. On the general strategy, see Shepard (2001) (and the commentaries published with that paper). That perceptual computation can draw on such general assumptions about environmental regularities goes some distance toward making up for the representational insufficiency of the featural array (see note 6) — it allows for ampliative expansion that would be otherwise unavailable.

(To be sure, it is possible to incorporate such regularities into a feature-prioritarian framework. However, we will argue in §2.2 that the cases we discuss are better accounted for by a no-priority conception of perception.)

that are relevant in determining their outputs, and also on the importance and variety of stored information they bring to bear on other inputs. For instance, the strategy for perceiving a unitarily hued color patch against a solid, non-contrasting background with no motion or depth cues will be simpler than the strategy for distinguishing objects in a more complicated visual field.

2.2 Examples: Color Perception, Depth Perception

We believe that the picture of perception as deploying, and then integrating the outputs of, a variety of perceptual strategies, is consonant with a broad range of contemporary empirical inquiry in perceptual science. As a way of sustaining this claim, let us consider a couple of representative examples.

A first example concerns color perception. Among the most pressing questions investigators have asked about color perception is how the perceptual system starts from the complex total signal striking the transducers — a signal that is determined jointly by the features of perceived objects and perceptual circumstances — and ends up with an assignment of colors to perceived objects.[19] How, for example, does the perceptual system start with the (spatially varying) array of light intensities reflected by the cup in figure 1 and end with the information that the entire cup is uniform in color (or, more cautiously, in some color-related respect)? Much of the research directed at answering this question has been aimed at developing algorithms for estimating the distinct contributions to the total signal made by the perceived object and the incident illumination: the hope is that, with these contributions factored apart, the system can subtract off (in Helmholtz's phrase "discount") the contribution of the illuminant in order to arrive at that portion of the signal due to the perceived object alone.

One of the things that is striking about this research is the impressively wide variety of distinct techniques for performing this factorization and arriving at representations of object colors (for an overview, see Cohen (2012)). For example, Maloney (1986); Maloney and Wandell (1986) show how a system with more classes of receptors than there are degrees of freedom in (the system's linear models of) surface reflection profiles can exploit its multiple receptoral signals to recover illumination-independent representations of surface features. Buchsbaum (1980) proposes a completely different model that rests on the assumption that the median lightness value in a scene corresponds to a middle grey surface, and computes from this assumption what the incident illumination would have to be to result in the observed intensity array. A related but distinct strategy proceeds from an assumption that anchors some

[19]We don't say this is the only or the most important use of color within the visual system; on the contrary, we agree with Akins and Hahn (2011); Kingdom and Mullen (1995); Kingdom (2003); Gegenfurtner (2003) that spectral information plays a variety of (integrative) roles in vision, of which assignment of colors to surfaces (and the like) is only one example. We take this point to be consonant with the general themes we are urging, and also a further reason to deny that color could play the role as a perceptual simple that it was assigned by the feature prioritarians discussed in §1.1.

part of the visual image to an extremal lightness value — for example, by treating the lightest visible surface as white (Land and McCann, 1971; Gilchrist *et al.*, 1999). (The latter approaches are further examples of strategies that take as inputs assumptions about the kinds of scenes perceptual systems will encounter.) Others have proposed estimating illuminants from information about mutual reflections in the scene (Funt *et al.*, 1991), the boundaries of regions known to be specular reflections (D'Zmura and Lennie, 1986; Lee, 1986), and shadows (D'Zmura, 1992). Still others propose to solve the problem by appeal to higher-order scene statistics, such as the correlation between redness and luminance within the scene (Golz and MacLeod, 2002) or the statistical distribution of colors within the scene (MacLeod, 2003; Brainard *et al.*, 2006). And, of course, as we discussed in §1.3.1, there is reason to think the relative contributions of surfaces and incident illumination are often easier to estimate if we start with uniformity assumptions about differently illuminated surface regions (i.e., if we take advantage of working models of object locations).

On its face, the situation manifestly appears to be one in which researchers have uncovered multiple perceptual strategies for the computation of surface color. Now, in principle it could turn out that there's just one true strategy that the perceptual system uses to the exclusion of all of the others. But we think this is extremely unlikely. For one thing, the different strategies listed above are differentially applicable to different scenarios: for example, the strategies that invoke known specular reflections, shadows, or object boundaries are only useful in cases where the perceptual system has access to those kinds of information. For another, no one of the proposals on the table so far correctly predicts the performance of the human perceptual system in all cases. Given this, it seems to us far more plausible that the perceptual system employs a variety of techniques (possibly even including some that are described above), and the combines their outputs to arrive at an integrative best guess about the way the world is.

A further example with the same moral involves depth perception. Here, too, there appear to be a variety of perceptual strategies by which the perceptual system assigns depth values to regions in the scene (some of which depend on objects). For example, investigators have proposed that, depending on the situation, depth can be computed from oculomotor cues such as accommodation (Wallach and Floor, 1971) and convergence (von Hosten, 1976); binocular disparity (Marr and Poggio, 1977, 1979; Nakayama and Shimojo, 1990; Jones and Malik, 1992); dynamic cues such as motion parallax (Rogers and Graham, 1979) and "optic flow" (Longuet-Higgins and Prazdny, 1980; Prazdny, 1980); pictorial cues such as line convergence, relative position, memory for known object types (Ittelson, 1951; Epstein, 1965), changes in regular texture gradients (Stevens, 1979; Kender, 1979; Witkin, 1981; Malik and Rosenholtz, 1994), and the interpretation of shading (Horn, 1975, 1977; Pentland, 1989; Kersten *et al.*, 1996) and edges (Guzman, 1968, 1969; Huffman, 1971; Clowes, 1971; Waltz, 1975; Malik, 1987).[20] Thus, once again, investigation has revealed a variety of perceptual strategies that plausibly contribute to the perception of depth. And, while in principle it could turn out that there is a uniquely correct and fully general perceptual strategy — that all of the others turn out not to be involved in depth perception after all, this appears unlikely. After all, as we found when considering color perception, the different perceptual strategies that appear to contribute to depth perception differ in their conditions of application, and are differentially successful in accounting for human performance over a broad range of conditions. What all this suggests is that, as in the case of color perception, depth perception proceeds by accumulating evidence from an assortment of perceptual strategies, and then integrates their outputs into a single representation.

2.3 Integration

We have said that the perceptual system deploys multiple perceptual strategies for computing the value of particular features such as surface color or spatial depth. If it is to deliver a univocal verdict about the world, then, perception must manage to integrate the outputs of these strategies in some way. Presumably the goal of such integrations is to arrive at a value that is consistent with (or, as consistent as possible with) the range of evidence supplied by the different perceptual strategies at work. Of course, different strategies may deliver inconsistent verdicts — particularly (but not exclusively) in laboratory settings designed specifically to pull them apart. And when this happens, the perceptual system must have some way of managing the conflicts, either by selectively preferring one strategy over others (e.g., as in our reaction to the Ames room, where perspective information about depth dominates over our expectations about the relative size of objects), by combining the information in some more complex way, by sending one or more of the strategies back to run on the input again in the hope of obtaining new answers that won't conflict, or by giving up (as, for example, we do with impossible Escher figures).

Similarly, the perceptual system must integrate the outputs of perceptual strategies operating over distinct spatiotemporal scales. Of course, the system is attempting to make global sense of the world, but does so on the basis of spatiotemporally local representations. It can happen, then, that a set of local assignments are all internally consistent, but fail to cohere globally. Once again, we expect that such clashes will require either an exclusive choice (dominance of one output over another), a complex combination, or giving up (possibly after a period of aspect switching between the locally coherent alternatives).

But if it is true that representations of different features — and indeed object representations — can all interact, then we should expect similar integration taking place between features, and between objects and features. As the cases discussed in §1.2 show, this is exactly what we find.

[20]See Palmer (1999, chapter 5) for a useful overview.

2.4 Perceptual Strategies and Modularity

On its face, the integrative conception of perception we are advocating might seem to be strongly at odds with the idea that perception is modular, or informationally encapsulated: the view's heavy reliance on interaction of different types of information as a basic functional element of perception seems to run counter to the basic idea of informational encapsulation.[21] However, we believe that the appearance of conflict between the integrative view and modularity is misleading, and that the former is at least compatible with, and perhaps supportive of, a modularist view.

To see why this is so, notice that the very notion of encapsulation (hence, modularity) presupposes boundaries. Only relative to a drawing of boundaries between a system and its surroundings is it meaningful to ask whether an effect is due purely to the internal functioning of the system. Now, feature prioritarianism is naturally connected with a specific view about how such boundaries should be drawn — viz., that they should sort representations (and the systems that process them) in terms of the features that are their contents. Given such feature-defined boundaries, the finding that perceptual representations of features interact with perceptual representations of objects (as in the cases detailed above) does constitute a transmission of information across a modular bound, hence a failure of informational encapsulation (hence, of modularity).

But why draw the boundaries this way? Unless we are already committed to feature-prioritarianism (as, we have argued, we should not be), it is more natural to construe the boundaries between modules in terms of the input that stereotypically mediates particular kinds of computations. Thus, for example, suppose that the specific apparent motion phenomena we've discussed are indicative of a perceptual strategy that computes the layout of objects and features over patterns of temporally extended stimulation (see §1.2.1). Then those patterns of temporally extended stimulation that mediate this response constitute the proprietary input for this perceptual strategy.

It could be objected that this trivializes the notion of modularity by introducing a disjunction — viz., a range — of input types. For instance, Prinz (2006, 28) argues that, trivially, "we can show that any collection of rules or representations is dedicated by simply listing an exhaustive disjunction of everything those rules and representations do". However, whether a particular set of inputs is disjunctive or not depends on the basic input vocabulary one uses, and this is exactly what is in dispute. According to a vocabulary of basic features, the current view is disjunctive; but according to a view based on perceptual strategies, it is not.[22] We suspect that the motivation for the trivialization worry stems from the concern that there should be principled

[21]For examples of this line of thought in recent writing on modularity, see Carruthers (2006); Prinz (2006); Stokes and Bergeron (2005).

[22]Indeed, we would point out that one can no less trivially save *anti*-modularity by fine-graining the taxonomy relative to the features Prinz has in mind — *round* is a disjunctive kind relative to a basic vocabulary whose primitives are *round-and-red* and *round-and-triangular*.

ways to distinguish the modular from the non-modular. In this case, the worry would be that whenever one finds an integration of a new type, one simply draws the boundaries around a new strategy, and therefore makes that integration part of a modular system. Thus, trivially, no system would ever be non-modular.

We can see that this objection fails, however, by considering paradigmatic cases of non-modularity, such as those involving ordinary practical reasoning. Suppose that one wants to go to the airport in the morning. Famously, the action planning that this desire causes needs to be responsive to an apparently unlimited variety of information, stemming from a variety of perceptual and cognitive sources. For instance, one's actions must be modified if one sees (via one's alarm clock) that one has overslept. But not if one remembers that one had mis-set one's alarm clock. Or, again, they could be influenced by remembering that one's flight is actually on Tuesday, or by being told that a hurricane has forced the airport to close, etc. The intended moral of these considerations is that neither the range of informational types to which the mechanisms responsible for producing this instance of practical reasoning are sensitive nor the range of operations defined over those informational types appear to be delimited in a way that would permit a systematic and sufficiently general theory of their operation. The system is too domain general to countenance a modular interpretation.[23] Putting aside the question of whether these problems about practical reasoning/central cognition are tractable, our point is that the integrative view on offer can indeed see a distinction between integrative processes that are potentially conducive to modularity (such as, perhaps, those discussed above) and those that are not (such as practical reasoning). This is enough to show that the trivialization objection under consideration fails.[24]

The lesson that emerges, then, is that the integrative view allows for modularity if and to the extent that the systems it targets operate on stereotyped inputs and perform stereotyped operations. Do these conditions obtain for some or many perceptual processes? As far as we can see, the integrative view is just agnostic on this (obviously empirical) matter. All we have attempted to establish here is that certain parade examples of what we take to be perceptual strategies are compatible with classical notions of modularity, despite their heavy reliance on integration.

Still, there is a worry that a certain sort of circularity is at work here. The boundaries that we've insisted must be in place before assessing modularity, and the computational properties of perceptual strategies, seem to be interdefined on our view. Surely, however, the integrative view is in no different a

[23]If these sorts of considerations sound familiar, they should: this sort of "isotropy" of central cognition was one of the factors that led Fodor (1983) to despair of a general, computational theory of mentation.

[24]Obviously, this distinction won't succeed if the perceptual integrations we are concerned with also turn out to draw on central states; for if so, then the non-encapsulation of the latter would threaten the modularity of the former as well. Crucially, however, nothing we've said about these strategies or their integration commits to their being connected to central states. If we are right about this, then the distinction we are drawing stands.

position in this regard than feature prioritarian views. It is true for everyone, as we've emphasized, that the notion of modularity is defined relative to a way of carving up the systems under study. Presumably, the best motivation for any particular carving would be that its divisions are enshrined in our best scientific theories and results about the domain. Our argument in the paper has been that, on the most comprehensive view of the evidence, the carving best suited to the study of perceptual systems is given in terms of perceptual strategies and their integration, rather than in the terms of feature prioritarianism. We claim that, relative to that taxonomy, there's no in principle conflict with modularity.

3 Conclusion

We have argued that extremely widespread and influential feature prioritarian views of perception are false. A considerable range of perceptual effects show that neither of the feature prioritarian's core commitments — to the representational sufficiency of features and the asymmetric dependence of object representations on feature representations — are generally true of perceptual processing. In turn, we have proposed an alternative account, based on perceptual strategies for information integration, that we claim both captures the effects we've presented and interacts fruitfully with a variety of productive projects in perceptual science. We suggest that the integrative view we've proposed provides a deeply different way of viewing both perceptual architecture and the mind's fundamental ways of connecting to the world, and has important implications for a number of empirical and conceptual issues about perception. We hope that the integrative account, or something like it, will prove fruitful in assessing philosophical questions about perception in ways consonant with the details of perceptual processing and with empirical advancements in the relevant domains.[25]

References

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga, editor, *The New Cognitive Neurosciences, 2nd Ed.*, pages 339–351. MIT Press, Cambridge, Massachusetts.

Akins, K. and Hahn, M. (2011). More than mere colouring. Ms., Simon Fraser University.

Bartels, A. and Zeki, S. (1998). The theory of multistage integration in the visual brain. *Proceedings of the Royal Society, B*, **265**, 2327–2332.

[25]Many thanks to Ned Block, Robert Briscoe, Craig Callender, Christopher Hill, Mark Kalderon, Don MacLeod, Mohan Matthen, Nicoletta Orlandi, Sam Rickless, Wayne Wu, and Chris Wüthrich for discussion and comments on earlier drafts, and to an audience at the American Philosophical Association, Pacific Division, who heard, and raised fruitful questions about, a version of this material.

Biederman, I. (1987). Recognition-by-components. *Psychological Review*, **94**, 115–147.

Blaser, E., Pylyshyn, Z. W., and Holcombe, A. O. (2000). Tracking an object through feature space. *Nature*, **408**, 196–199.

Braddick, O. (1974). A short-range process in apparent motion. *Vision Research*, **14**, 519–527.

Brainard, D. H., Delahunt, P. B., Freeman, W. T., Kraft, J. M., and Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, **6**, 1267–1281.

Buchsbaum, G. A. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, **310**, 1–26.

Carnap, R. (1967). *The Logical Structure of the World; Psuedoproblems in Philosophy*. University of California Press, Berkeley, California. Translated by Rolf A. George.

Carruthers, P. (2006). The case for massively modular models of mind. In R. Stainton, editor, *Contemporary Debates in Cognitive Science*, pages 3–21. Blackwell, Malden, MA.

Clark, A. (2004). Feature placing and proto-objects. *Philosophical Psychology*, **17**(4), 443–469.

Clowes, M. (1971). On seeing things. *Artificial Intelligence*, **2**(1), 79–116.

Cohen, J. (2004). Objects, places, and perception. *Philosophical Psychology*, **17**(4), 471–495.

Cohen, J. (2008). Color constancy as counterfactual. *Australasian Journal of Philosophy*, **86**(1), 61–92.

Cohen, J. (2012). Perceptual constancy. In M. Matthen, editor, *Oxford Handbook of Philosophy of Perception*. Oxford University Press, Oxford. forthcoming.

Cornsweet, T. (1970). *Visual Perception*. Academic Press, New York.

Craik, K. J. W. (1966). *The Nature of Psychology*. Cambridge University Press, Cambridge.

D'Zmura, M. (1992). Color constancy: surface color from changing illumination. *Journal of the Optical Society of America A*, **9**, 490–493.

D'Zmura, M. and Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, **3**, 1662–1672.

Epshtein, B. and Ullman, S. (2005). Feature hierarchies for object classification. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 - Volume 01*, pages 220–227, Washington, DC, USA. IEEE Computer Society.

Epstein, W. (1965). The known-size apparent-distance hypothesis. *American Journal of Psyschology*, **74**, 333–346.

Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press, Cambridge, Massachusetts.

Funt, B., Drew, M., and Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, **6**, 5–24.

Gegenfurtner, K. R. (2003). Cortical mechanisms of colour vision. *Nature Reviews Neuroscience*, **4**, 563–572.

Gilchrist, A. L., Kossifydis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., and Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, **106**(4), 795834.

Golz, J. and MacLeod, D. I. A. (2002). Influence of scene statistics on colour constancy. *Nature*, **415**, 637–640.

Goodman, N. (1951). *The Structure of Appearance*. Harvard University Press, Cambridge, Massachusetts.

Guzman, A. (1968). *Computer Recognition of Three-Dimensional Objects in a Visual Scene*. Ph.D. thesis, MIT.

Guzman, A. (1969). Decomposition of a visual scene into bodies. In A. Griselli, editor, *Automatic interpretation and classification of images*, pages 243–276. Academic Press, New York.

Horn, B. (1975). Obtaining shape from shading information. In P. H. Winston, editor, *The psychology of computer vision*, pages 115–155. McGraw-Gill, New York.

Horn, B. (1977). Understanding image intensities. *Artificial Intelligence*, **8**(2), 201–231.

von Hosten, C. (1976). The role of convergence in visual space perception. *Vision Research*, **16**(2), 139–144.

Huffman, D. A. (1971). Impossible objects as non-sense sentences. In M. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 6. Edinburgh University Press, Edinburgh.

Ittelson, W. H. (1951). Size as a cue to distance. *American Journal of Psychology*, **64**, 54–67.

Jones, D. G. and Malik, J. (1992). Computational framework for determining stereo correspondence from a set of linear spatial filters. *Image Vision Computation*, **10**, 699–708.

Kender, J. (1979). Shape from texture: A computational paradigm. In *Proceedings of the DARPA Image Understanding Workshop*, pages 134–138.

Kersten, D., Knill, D. C., Mamassian, P., and Bülthoff, I. (1996). Illusory motion from shadows. *Nature*, **279**(6560), 31.

Kingdom, F. A. A. (2003). Colour brings relief to human vision. *Nature Neuroscience*, **6**, 641–644.

Kingdom, F. A. A. and Mullen, K. T. (1995). Separating color and luminance information in the visual system. *Spatial Vision*, **9**, 191–219.

Koffka, K. (1935). *Principles of Gestalt Psychology*. Harcourt, Brace, & World, New York.

Kolers, P. and Grunau, M. V. (1976). Shape and color in apparent motion. *Vision Research*, **16**, 329–335.

Land, E. H. and McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, **61**.

Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, **3**, 1694–1699.

Locke, J. (1975). *An Essay Concerning Human Understanding (1689)*. Oxford University Press, New York.

Longuet-Higgins, H. C. and Prazdny, K. F. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society, London, B*, **208**, 385–397.

MacLeod, D. I. A. (2003). Colour discrimination, colour constancy, and natural scene statistics. In J. D. Mollon, J. Pokorny, and K. Knoblauch, editors, *Normal and Defective Colour Vision*. Oxford University Press, London.

Malik, J. (1987). Interpreting line drawings of curved objects. *International Journal of Computer Vision*, **1**(1), 73–103.

Malik, J. and Rosenholtz, R. (1994). A computational model for shape from texture. *Ciba Foundation Symposium*, **184**, 272282.

Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, **3**, 1673–1683.

Maloney, L. T. and Wandell, B. A. (1986). Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, **3**(1), 29–33.

Marr, D. and Poggio, T. (1977). Cooperative computation of stereo disparity. *Science*, **194**, 283–287.

Marr, D. and Poggio, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society of London, B*, **204**, 301–328.

Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford University Press, Oxford.

Nakayama, K. and Shimojo, S. (1990). DaVinci stereopsis: Depth and subjective contours from unpaired monocular points. *Vision Research*, **30**, 1811–1825.

O'Brien, V. (1959). Contrast by contour-enhancement. *American Journal of Psychology*, **72**, 299–300.

Orlandi, N. (2010). Are sensory properties represented in perceptual experience? *Philosophical Psychology*, **23**(6), 721–740.

Palmer, S. E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences*, **22**(6), 923–989.

Pentland, A. (1989). Shape information from shading: A theory about human perception. *Spatial Vision*, **4**(3/4), 165–182.

Prazdny, K. (1980). Egomotion and relative depth from optical flow. *Biological Cybernetics*, **36**, 87–102.

Prinz, J. J. (2006). Is the mind really modular? In R. Stainton, editor, *Contemporary Debates in Cognitive Science*, pages 22–36. Blackwell, Oxford.

Purves, D., Shimpi, A., and Lotto, R. B. (1999). An empirical explanation of the Cornsweet effect. *The Journal of Neuroscience*, **19**(19), 8542–8551.

Pylyshyn, Z. W. (2003). *Seeing and Visualizing: It's Not What You Think*. MIT Press, Cambridge, Massachusetts.

Pylyshyn, Z. W. (2004). Some puzzling findings in multiple object tracking (MOT): I. tracking without keeping track of object identities. *Visual Cognition*, **11**, 801–822.

Pylyshyn, Z. W. and Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, **3**(3), 1–19.

Quine, W. V. O. (1960). *Word and Object*. MIT Press, Cambridge, Massachusetts.

Riesenhuber, M. and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, **2**, 1019–1025.

Rogers, B. and Graham, M. (1979). Motion parallax as an independent cue for depth perception. *Perception*, **8**(2), 125–134.

Russell, B. (1914). *Our Knowledge of the External World*. Open Court, Chicago, IL.

Schiller, A. (2012). The primacy of fact perception. *Philosophical Psychology*. in press.

Scholl, B. J. and Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, **38**, 259–290.

Scholl, B. J., Pylyshyn, Z., and Feldman, J. (2001). What is a visual object? evidence from target merging in multiple object tracking. *Cognition*, **80**, 159–177.

Shepard, R. N. (2001). Perceptual-cognitive universals as reections of the world. *Behavioral and Brain Sciences*, **24**, 581–601.

Shimojo, S., Scheier, C., Nijhawan, R., Shams, L., Kamitani, Y., and Wantanabe, K. (2001). Beyond perceptual modality: Auditory effects on visual perception. *Acoustical Science & Technology*, **22**, 61–67.

Siegel, S. (2006). Which properties are represented in perception? In T. S. Gendler and J. Hawthorne, editors, *Perceptual Experience*, pages 481–503. Oxford University Press, Oxford.

Spelke, E. (1990). Principles of object perception. *Cognitive Science*, **14**, 29–56.

Steinman, R. M., Pizlo, Z., and Pizlo, F. J. (2000). Phi is not beta, and why Wertheimer's discovery launched the Gestalt revolution: a minireview. *Vision Research*, **40**, 2257–2264.

Stevens, K. A. (1979). *Surface perception from local analysis of texture and contour*. Ph.D. thesis, MIT.

Stokes, D. and Bergeron, V. (2005). Modular architectures and informational encapsulation: A dilemma. Unpublished ms.

Treisman, A. (1998). Feature binding, attention and object perception. *Phil Trans R. Soc London B*, **353**, 1295–1306.

Ullman, S. (1979). *The Interpretation of Visual Motion*. MIT Press, Cambridge, Massachusetts.

Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences*, **11**, 58–64.

Wallach, H. and Floor, L. (1971). The use of size matching to demonstrate the effectiveness of accomodation and convergence as cues for distance. *Perception & Psychophysics*, **10**, 423–428.

Waltz, D. (1975). Understanding line drawings of scenes with shadows. In P. H. Winston, editor, *The psychology of computer vision*, pages 19–92. McGraw-Gill, New York.

Westphal, G. and Würtz, R. P. (2009). Combining feature-and correspondence-based methods for visual object recognition. *Neural Computation*, **21**, 1952–1989.

Witkin, A. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence*, **17**(1-3), 17–45.

Yuille, A. and Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, **10**, 302–308.

Zeki, S. (2001). Localization and globalization in conscious vision. *Annual Review of Neuroscience*, **24**, 57–86.