



PROJECT MUSE®

---

## Is Emotion a Form of Perception?

Prinz, Jesse J.

Canadian Journal of Philosophy, Volume 36 (2006) Supplement [vol. 32], pp. 137-160 (Article)

Published by Canadian Journal of Philosophy  
DOI: [10.1353/cjp.2007.0035](https://doi.org/10.1353/cjp.2007.0035)



➔ For additional information about this article

<http://muse.jhu.edu/journals/cjp/summary/v036/36.5Sprinz.html>

## *Is Emotion a Form of Perception?*

JESSE J. PRINZ

Theories of emotions traditionally divide into two categories. According to some researchers, emotions are or essentially involve evaluative thoughts or judgments. These are called *cognitive theories*. According to other researchers, an emotion can occur without any thought. These are called *non-cognitive theories*. Some defenders of non-cognitive theories argue that emotions are action tendencies, others say they are feelings, and still others say they are affect programs, which encompass a range of internal and external events. One of the most celebrated non-cognitive theories owes, independently, to William James and Carl Lange. According to them, emotions are perceptions of patterned changes in the body. I think the perceptual theory of emotions is basically correct, but it needs to be updated. In this discussion, I will offer a summary and defence.

The question I am addressing bears on the question of modularity. Within cognitive science, there is a widespread view that perceptual systems are modular. If this is right, then showing that emotion is a form of perception requires showing that emotion is a modular process, and showing that emotion is modular could contribute to showing that emotion is a form of perception (assuming that not all mental capacities are underwritten by modular systems). Therefore, modularity will figure centrally in the discussion that follows, as it did in an earlier treatment of this topic (Prinz 2004). There is, however, a change in how I will approach this topic here. I have come to believe that perception is not, in fact, modular as that term is defined by Fodor (1983) in his classic treatment of the topic. Perceptual systems bear features in common with Fodor's modules, but Fodor's approach is, in my view, mistaken (Prinz 2006). Here I will introduce the idea of quasi-modules, which bear some things in common with Fodor's modules,

and I will argue that emotions are quasi-modular. This thesis will help secure the parallel between emotion and perception.

## I. What Is Perception?

To determine whether emotion is a form of perception, it would be handy to have a working definition of perception. I am not going to offer such a definition, however. Offering conceptual analyses of psychological terms is methodologically unscrupulous. It presupposes, quite implausibly, that our ordinary folk psychological terms have good definitions. There are some paradigm cases of perception, but what these have in common must be determined by careful observation and theory construction, not armchair lexicography. Still, as a starting place, we can reflect on some of the features that paradigm perceptions have in common. If we consider visual, auditory, and olfactory states, for example, we find the following characteristic features:

First, perception takes place in sensory systems. Sensory systems are systems that convert physical magnitudes into mental representations. Each sensory system has dedicated transducers that are stimulated by non-mental features of the world, and output mental representations in a modality-specific code.

Second, perception involves the generation of internal representations, and these typically represent the mind-external stimuli. Sometimes the senses represent proximal stimuli (i.e., perturbations of our sensory transducers), but they can also represent more distal stimuli (e.g., external objects) or relational properties (e.g., secondary qualities, or powers that external objects have to cause mental states in us). It is important to emphasize that sensory systems may have to do a fair amount of processing before representations of complex distal objects can be generated. When you see a giraffe, for example, the eyes first convert light reflected from the surface of the giraffe into a vast assembly of edge representations and colour patches. These are then bound together and organized into a representation of the giraffe's contours. Those contours are used to extract perceptual invariants that remain constant across various viewing positions and these are matched against stored templates in visual memory. Through this process, the visual system ends up generating a giraffe representa-

tion. It generates that representation by first representing something else: patterns of light. Giraffes are not identical to their appearances, but we detect them through their appearances. Objects are not directly given to the senses. They must be reconstructed or extrapolated from the superficial magnitudes that senses transduce.

Third, perceptions can be consciously experienced. They have phenomenal qualities. This is significant because, arguably, perceptual representations are the only internal states of which we can be conscious. When we consider our phenomenal qualities, all of them seem to be modality-specific. We can recognize a smooth surface with eyes or touch, but smoothness is not presented consciously in an amodal code; it always presents itself in consciousness as visual or tactile. Even thoughts present themselves to us in modality-specific codes. We experience thoughts as images of the facts they represent or, more commonly, as strings of words in the languages we speak. When thinking about philosophy, for example, we usually hear auditory images of sentences running through our heads. There is no uncontroversial example of a phenomenal quality that is not perceptual in character. I will assume throughout that only perceptual states are phenomenally conscious, and I will also suggest below that perceptual states become phenomenally conscious in exactly the same way.

Fourth, perception is quasi-modular. I add the ugly prefix because, as I remarked at the outset, I do not think perception is modular in the way that the term is defined in Fodor's (1983) influential book. I have argued against Fodorian modularity at length elsewhere (Prinz 2006). Rather than rehearsing those arguments, I will be a bit more constructive here. For even if perceptual systems are not modular in the standard sense that has been given to the term, they share some features in common with modules. I will describe these features and label any system that has them *quasi-modular*.

A mental capacity is *quasi-modular* to the extent that it is:

1. Functionally specialized
2. Subject to characteristic breakdowns
3. Capable of automatic processing
4. Built up from a system of innate rules and representations
5. Stimulus-dependent

On the face of it, all of these criteria are consistent with Fodor's definition of modularity. He talks about domain specificity, characteristic breakdowns, automatic processing, and innateness. But the last item on my list, stimulus-dependence, is introduced to replace the aspect of his definition that he considers most fundamental to modularity: informational encapsulation. By stimulus-dependent, I mean to suggest that modules are constrained by their inputs. Top-down influences cannot fully determine what happens inside a module. In terms of perception, the idea is that we cannot simply choose what we perceive. But stimulus-dependence is consistent with the possibility that top-down influences can significantly affect what happens in perceptual systems. And this is an important departure from Fodor's idea of "encapsulation." Fodor insists that perceptual systems do not let in any information from systems further up the information-processing hierarchy, and I think he is mistaken about this (see Prinz 2006). There is overwhelming evidence that cognitive systems can communicate with perceptual systems (as in the case of mental imagery), and that perceptual systems can communicate with each other (as with the McGurk effect [McGurk and MacDonald 1976] and other forms of intermodal accommodation).

Fodor argues for informational encapsulation by appeal to perceptual illusions. Consider the Müller-Lyer illusion, in which equal lines appear different in length. This illusion persists even after we learn that the lines are the same, which leads Fodor to conclude that knowledge cannot penetrate the perception system. But this conclusion is too strong. An alternative explanation is that bottom-up inputs *trump* top-down inputs when the two come into conflict. After all, when there is no conflict, there can be top-down influences. Consider the duck-rabbit. Cognition can lead us to reconstrue the image as a duck after seeing it as a rabbit, because both interpretations are consistent with the stimulus. I think we should drop talk of encapsulation and appeal to the notion of trumping when describing perceptual systems. The notion of stimulus-dependence captures that idea. In saying that we cannot choose what we perceive, I mean to imply that we cannot, under ordinary circumstances, have a perceptual representation generated top-down when the current stimulus is disposed to induce an incongruent representation. Perceptual systems are stimulus dependent in this sense, and they are quasi-modular.

In sum, while I have not offered a definition of perception, I have identified four features that are characteristic of perception. Perception takes place in modality-specific input systems; perception involves the generation of representations, which represent proximal, distal, or relational stimuli; third, perceptions can be consciously experienced; and fourth, perceptions are quasi-modular. In arguing that emotion is a form of perception, I will consider each of these features in turn. If I can establish that emotions exhibit all four features, then, to that extent, emotions are like paradigm instances of perception. I don't want to insist that each of these features is necessary for qualifying as a case of perception, but collectively they strike me as a plausible set of sufficient conditions. No one condition is sufficient on its own to establish that emotions are perceptions. Perhaps a proper subset of these conditions would suffice, but I will remain neutral about that question here. I will argue that emotions exhibit all four features, and thus deserve to be called perceptions. This is not intended as a conceptual claim. I mean to argue that it is an empirical fact that emotions share features that can be empirically observed in paradigm instances of perception. Getting mad is very much like seeing red.

## **II. Emotions Are Modality-Specific**

I think emotions qualify as states of modality-specific input systems. They are interoceptive; emotions are states in the sensory systems that respond to changes in the body. I will defend the interoception thesis in this section. To remain neutral about the question of whether emotions are *perceptions* of the body, as opposed to merely being states that happen to occur in a sensory system, I will adopt a distinction between registration and representation. I will say that a response within a perceptual system that reliably occurs in response to a stimulus *registers* that stimulus. To *represent* rather than merely registering, the state would have to have the function of being caused by the stimulus. A mental state represents that which it has the function of reliably detecting (cf. Dretske 1988). Correlatively, to qualify as a perception of a stimulus, the state would have to represent it. States of sensory systems that register stimuli can be called "perceptual states" simply in virtue of being modality-specific, but to grant them the title "perceptions" we might demand that they also represent the stimuli

that cause them. It will be the burden of this section to show that emotions are perceptual states, but the burden of the next section to show that they represent.

The link between emotions and interoception was most influentially explored by William James (1884) and Carl Lange (1885). On more traditional folk psychological theories, emotions were presumed to cause various bodily changes, such as crying or trembling or fleeing. For James and Lange, this gets things back to front. Emotions are the effects of bodily changes, not the causes. Crying makes us sad, or, more accurately, the feeling of sadness is the feeling of a range of bodily responses that includes crying. On this view, emotions arise in the following way. We encounter a stimulus or have a thought that is (as the result of evolution or learning) disposed to trigger a pattern of changes in our bodies; when those changes occur, they are registered in brain systems that are sensitive to somatic states; these neuronal responses are experienced as feelings, and those feelings are what we call emotions. Lange focuses on vascular changes, but James, following Darwin, argues that each emotion is associated with a complex range of bodily changes and emotions feel different as a result of the different bodily patterns that they register. When frightened, we tremble, and breathing becomes strained. When angry, our muscles tighten, our brows lower, and blood rushes to our face and extremities. When overjoyed, our hearts race, our breathing becomes relaxed, and our arms widen receptively. Our circulatory systems, respiratory systems, facial expressions, and bodily movements are all correlated with emotions. When inferring causation from correlation, we tend to think that emotions are the causes of these varied bodily changes, rather than the effects, but James reverses the order of events. The brain contains interoceptive systems that are linked to the body by a vast network of nerve fibres. Like the rods and cones in our eyes, these nerves serve as transducers converting physical changes into electrochemical signals that are sent into the brain. The interoceptive systems that register patterned changes in the body are the neural substrates of emotions on the James-Lange view. Emotions are interoceptive states that register patterned changes in the body. Call this the registration thesis.

Both James and Lange defend the registration thesis by means of a thought experiment. They ask us to imagine having an emotion without having any changes in the body. Imagine terror, for example, with-

out any spine-tingling, trembling, or constrained breathing. Imagine that your body is completely relaxed. James and Lange think it will be obvious that there is nothing in this placid bodily state that could be recognized as terror. Subtract the bodily perturbation, and the emotion goes as well. This suggests that emotions are interoceptive states.

This argument rests on introspective intuitions, which may not be shared by everyone. Lange comments that it is easy to deceive yourself into thinking that you can have an intense emotional experience in the absence of any somatic feelings. Somatic feelings are often not recognized as such. When we have a twinge or muscle tension or a subtle imbalance in our vestibular systems, we do not always recognize the cause. So the appeal to introspection will not persuade everyone that James and Lange are right. We need other sources of evidence. In recent years, there has been a growing body of empirical evidence in support of the registration thesis (for reviews, see Damasio 1994; Prinz 2004). Here I mention just a few lines of evidence.

First, there is now a large body of evidence suggesting that emotions can be induced by changing states of the body. One method is to use drugs that act on the autonomic nervous system. When people are injected with adrenalin, they report having experiences that feel like emotions (Marañón 1924). Another method uses feedback from the body. When we make facial expressions (Strack et al. 1988), or change our posture (Stepper and Strack 1993), or change our breathing pattern (Philippot et al. 2002), we often experience a corresponding emotion. Smile and you will feel happy; scowl and you will feel mad. This is just what the registration thesis predicts.

Further evidence for the registration thesis comes from neuroimaging. There have now been hundreds of studies of brain activity during emotion episodes. Again and again, these studies implicate the same structures (Phan et al. 2004; Wiens 2005). The structures that get discussed most frequently are the amygdala, which is primarily believed to be involved in the induction of emotions, and the cingulate and insular cortices, which are implicated in the emotions themselves. The amygdala is essentially an association area that links perceptions or thoughts to bodily responses. When an emotionally significant event is perceived or contemplated, the amygdala sends signs to brain structures that regulate changes in the endocrine system, the autonomic nervous system, and in systems that control stereotyped behav-



ournal responses (LeDoux 1996). Both the cingulate and the insula are involved in the regulation and response to changes in the body. Both of these complex brain areas are implicated in interoception (Critchley et al. 2004; Wiens 2005). The fact that emotions reliably correlate with activity in these brain areas, and do not systematically correlate with activity in other areas, provides support for James and Lange's suggestion that emotions are interoceptive states.

The evidence from bodily feedback suggests that bodily changes can be sufficient for emotions, and the evidence from neuroimaging suggests that emotions co-occur with, and may ordinarily be constituted by states in brain systems that register bodily changes. This is strong evidence in favour of the registration thesis that was expounded by James and Lange. But critics of this tradition may complain that, while perceptions of the body are sufficient for emotions, they are not necessary. They may argue that some emotions arise in the absence of bodily changes and perceptions thereof. In response, defenders of the registration thesis could either concede that some emotions are disembodied, as it were, or they could dig in their heels and argue that all emotions are perceptions of bodily states. I think a qualified version of heel-digging is defensible. Let me introduce two qualifications:

First, emotions can be attributed as states or as traits. As an example of an emotion trait, consider the sentence, "Scottie is afraid of heights." This statement is true of Scottie even when Scottie isn't experiencing the state of fear. Emotion traits are not perceptions of bodily changes. But, I think a person can be truly attributed an emotion trait only if he or she is disposed to have the corresponding state (Scottie is afraid of heights if and only if heights instill fear in him), and emotion states, I claim, are always perceptions of bodily changes. I defend this claim in more detail elsewhere (Prinz 2004).

Second, emotions sometimes seem to arise before we've had time to perceive bodily changes (Cannon 1927), and they can even arise in individuals who have limited capacity to perceive bodily changes due to spinal cord injuries (Chwalisz et al. 1988). These points are usually presented as an objection to the James-Lange approach, but, as Damasio (1999) points out, there is an easy response. For James and Lange, emotions are interoceptive states – states in brain systems that normally register changes in the body. Presumably such states

can arise even in the absence of actual bodily changes. Just as we can visually imagine an object without seeing it, we can imagine bodily changes without those changes taking place. It is extremely plausible that, when we encounter familiar emotion elicitors, our interoceptive systems become active before the body has had time to change. The brain anticipates what the body will do. Thus, the claim that emotions are interoceptive states is compatible with the claim that emotions can occur before bodily changes and when bodily changes are imperceptible due to spinal injury. This may appear *ad hoc*, but it's not. There is evidence that when bodily changes are executed endogenously, as when we move our limbs, motor and somatosensory systems generate a "forward model" that predicts what the executed bodily changes will be like (Wolpert and Flanagan 2001). When acting, we generate images of our actions before they have even occurred. The present suggestion is that interoceptive systems use forward models as well. This proposal requires a slight refinement of the registration thesis. Emotions, I will say, are states of interoceptive systems that either register or *anticipate* changes in the body. In cases where emotions involve interoceptive states that merely anticipate bodily changes rather than registering them, emotions are more like perceptual *images* than perceptions. But this qualification does not vitiate the claim that emotion is, in general, a form of perception. If it did, we would have to deny that vision is a form of perception. After all, we can form visual imagery, and we often generate visual images in anticipation of the objects that we are about to see (Kosslyn 1994). When watching a moving ball, we may spontaneously use imagery to anticipate its trajectory; likewise, when responding to an impending threat, we may spontaneously anticipate how our bodies will feel once the fear response has been fully activated.

In sum, I think a strong case can be made for James and Lange's proposal that emotions are interoceptive states that (normally) serve to register changes in the body. The registration thesis is consistent with evidence from psychology and neuroscience. The research on bodily feedback suggests that body changes can cause emotions, and neuroimaging results show that emotions supervene on brain structures that register changes in the body. Taking these findings at face value, we should conclude that emotions are interoceptive states.

Given the evidence, opponents of this view have the burden of proof. I am convinced that James and Lange were right.

If emotions are interoceptive states, then emotions are perceptual states because interoceptive systems are perceptual systems; interoception is one of our senses. But are emotions perceptions? To qualify as perceptions, emotions must be perceptions *of* something. Is having an emotion perceiving anything? If so, what?

### III. Emotions Are Representations

So far, I have been very cautious about terminology. I said that emotions are perceptual states, not that they are perceptions. In particular, I have not said that emotions allow us to perceive anything. To defend that claim, I must establish that emotions represent something. Perceptions are representations. On the face of it, this might look like a trivial step from the evidence that I have been presenting. After all, I have been arguing that emotions register patterned changes in the body. Can't I just switch terminology and say that emotions represent such changes? Aren't emotions perceptions of bodily changes? Well, perhaps not. In this section I want to argue that, while emotions register bodily changes, they don't represent them. Indeed, this conclusion is important because, if emotions were representations of bodily states, they couldn't be distinguished from itches, twinges, or chills. Emotions have semantic properties that distinguish them from garden variety bodily perceptions. Let me explain.

According to the theory of representation mentioned in section II, representation requires that two conditions be met. A mental state represents that which it has the function of reliably detecting (for a defence, see Dretske 1988; Prinz 2000). The reliable detection condition is causal. Under ordinary conditions, mental representations are causally activated when we contact their referents. When we see red, our red representations are tokened, and when we see water our water representations are tokened. But many different things can cause a mental state to activate. Hearing the word "rose" may cause a red image to activate, and a long trek in the desert may lead us to token a representation of water. If mental representations referred to all of their causes, red would refer to the word "rose" and our water concepts would refer to the arid desert. To avoid this unwanted consequence,

we should following leading theories in psychosemantics and say that mental representations represent only those of their many causes that they have the function of detecting. In this context, “function” can be cashed out in terms of causal history. We can ask, how was the mental representation in question acquired in the course of learning or natural selection. If a mental representation was passed on through natural selection as a result of our ancestors successfully detecting red things, then it represents red. If a mental representation was learned as a consequence of being presented with samples of water, then it is a water representation.

Assuming this psychosemantic theory is correct, we can ask, what do emotions have the function of detecting? I have already suggested that emotions detect patterned changes in the body. These things reliably cause emotions to occur. But do they have the function of detecting such changes? Let’s assume, for these purposes, that our emotions are the product of natural selection. I think this is true of some emotions and others are learned by deploying evolved emotions in new contexts. Focusing on evolved emotions, we can ask, did our capacity to detect patterned bodily changes get passed onto us in virtue of the fact that they co-varied with those bodily changes, or for some other reason? Like any question about our evolutionary history, this is difficult to answer, but I think we can safely speculate by considering the current function of emotions. Emotions are, if the registration theory is right, reliable indicators of body states. They carry the information that our hearts are beating and our lungs are contracting. The body does need to monitor such bodily states in order to regulate the basic biological functions necessary for life. But information about our viscera is not essential to the functional role of *emotions* in regulating behaviour. Emotions help us choose behavioural responses that cope with external situations, not with internal organs. But how do they do that?

The trick is to use the body as an indicator of how we are faring in the world. Think about how smoke alarms work. They are wired to emit a sound when smoke is near. Likewise, we are wired to enter into patterned bodily states when matters of concern arise. Evolution has set things up so that we enter into a distinctive bodily pattern when we encounter certain dangerous things (loud sudden noises, predators, sudden loss of support), another bodily pattern when we

encounter certain threatening things (a glare, theft, or attack from a conspecific), and a third pattern when we encounter certain losses (the death or disappearance of a conspecific with whom we have a close relationship). For each biologically basic emotion, there is a distinctive bodily pattern and a distinctive set of eliciting conditions. Each set of eliciting conditions instantiates a specific kind of organism-environment relation that bears on well-being (danger, threat, loss, and so on). Call these relations “concerns.” Because we are wired in this way, our bodily patterns *and the brain states that register them* are reliably caused by concerns. Quite plausibly, it is in virtue of detecting these concerns, and not the bodily states themselves, that our body-pattern detectors got passed down from our ancestors.

Think of it this way. There is good evolutionary reason why we should be able to register local changes in the body, such as a racing heart or constricted breath. Bodily homeostasis requires feedback from the organs regulated by the central nervous system. But why do we need to register *patterns* of bodily change? Why do we need brain states that register a racing heart *together with* strained breathing? Homeostasis may not require that. So it is an evolutionary puzzle why we would be able to detect patterns of bodily change. The puzzle is solved if we imagine that these patterns ordinarily occur in special circumstances. If the heart races together with strained breathing under situations that were dangerous to our ancestors, then the brain could use this information to register danger, and, more specifically, it could use that signal to tell action-selecting systems to search for response strategies that are useful for avoiding danger. The body, like the tone in the smoke alarm, signals that there is something we need to cope with in our environment.

What I am suggesting is that a neural response to a patterned bodily change causally covaries with two different things. The response covaries with the bodily change in question but also with a concern, such as danger, threat, or loss. The neural response detects both of these things, but it represents only the one that it has the function of detecting — the one that it was selected for detecting. It seems very plausible that neural responses to patterned bodily changes have the function of detecting concerns. Such neural responses register bodily changes, but they thereby represent concerns such as danger, threat, and loss.

This conclusion may seem to be at odds with the thesis that emotion is a kind of perception. It seems reasonable, on the face of it, to say that neural states in our interoceptive systems are perceptions of bodily changes, but bizarre to say that they are perceptions of danger, threat, and loss. Those concerns seem too abstract and too disparate in form to be perceived. After all, dangers don't share any morphological properties in common. They don't look alike. There is no obvious set of appearances uniting all and only dangers. In paradigm cases of perception, we assume that the perceived property is a superficial appearance and that the perceptual representation of that property resembles it, in some respect. We think that a visual perception of a circle consists in a perceptual representation that is circular. Emotions don't resemble the concerns that I have been considering, nor could they. Those concerns are instantiated by concrete objects and events that are too variable in form.

The worry is easily addressed once we realize that the resemblance theory of perception is false — even in paradigm cases. Assume, for example, that colours are secondary qualities. Assume that red is the power to cause a certain experience in us. Two things follow from this. First, the property of being red is morphologically heterogeneous. The physical entities that have the power to cause red experiences in us are highly varied. They have nothing intrinsic in common in virtue of which they might be grouped together. Their unity lies in the effect they have on us. Second, our perceptual representations of red do not resemble what they represent. Red is represented as a specific phenomenal quality. There is nothing out there that is intrinsically red. The world, without us, is colourless. So red experiences do not resemble what they represent. Some people deny that red is a secondary quality. I don't want to take a stand on that debate. The point is that perceptual representations can represent complex relational properties that have no intrinsic morphological unity, and they can represent without resembling their referents. If this is even a *possibility* in the case of red, we should not rule out the possibility that emotions represent abstract relational properties such as danger, threat, or loss.

I conclude that emotions do represent concerns. It will take a little more work, however, to show that they are *perceptions* of concerns. I will come to this point in the final section, but first a detour through consciousness.

## IV. Emotions Are Perceptually Conscious

The fact that emotions feel like something is a tip-off to the fact that they are perceptual. After all, every other phenomenally conscious state seems to be perceptual in nature. We have visual experiences, auditory experiences, olfactory experiences, tactile experiences, and so on. Beyond the experiences associated with each sense modality, there seem to be no phenomenal qualities. As remarked above, even conscious thoughts seem to come to us in the form of mental images of what those thoughts denote or images of the words we would use to express them. Subtract away all conscious perceptual representations and we have nothing left in consciousness at all. This is a controversial position, but I see no compelling example of a conscious quality that cannot be pinned on a particular sensory system. If emotions are conscious, they must be sensory.

But what are conscious emotions like? How should we describe their qualitative character? I have just been arguing that emotions represent abstract relational properties, such as danger and loss. Should we say that fear feels dangerous? In a certain sense, this way of talking is completely appropriate. Fear is a way of detecting danger, and consequently dangerous things cause fear in us. If danger feels like anything at all, it feels like fear. We can say, "I feel a sense of impending danger" or "that looks hazardous," meaning that some situation has induced in us a feeling of fear. But the feeling of fear can also be described in another way. If you attend to the phenomenology of fear, you will notice that the characteristic qualities of the experience are somatic. Muscles tighten, the body arches back or cowers, hairs stand on end, eyes widen, blood vessels become constricted, the heart races, and breath shortens. As James argued, these bodily changes seem to exhaust the feeling of fear.

If this story about the qualities of emotional experience is right, then emotional consciousness is a species of perceptual consciousness. This is a satisfying discovery because it lends itself to a unified theory of conscious experience. Most of our conscious episodes are *obviously* perceptual. At any given moment, we are seeing things, hearing things, touching things, and perhaps smelling and tasting things. Arguably, the sum total of conscious qualities at any given moment can be entirely explained in perceptual terms. Even the con-

scious experience of thinking can be characterized perceptually. Most typically, conscious thoughts are verbalized. As you mull over a philosophical hypothesis, you describe it to yourself in silent speech. The experience of thinking is the experience of hearing sentences in your head. Of course, thoughts are also sometimes accompanied by other forms of imagery; we may visualize the things that sentences describe. All this, I submit, can be explained in sensory terms (Prinz, 2007). But, in addition to silent speech and visual images, our conscious thoughts and experiences are also often accompanied by conscious emotions. If emotional consciousness could not be explained in perceptual terms, we would be stuck with the odd conclusion that all consciousness is perceptual except emotional consciousness. This would be an inelegant outcome. All conscious states share something in common, namely the fact that they have phenomenal qualities. It would be nice to have a unified theory of phenomenal qualities. If all consciousness were perceptual, a unified theory might be possible, but, if emotional consciousness were not a form of perceptual consciousness, then we might have to develop two theories of consciousness – one for perception and the other for emotion. Fortunately, I think emotional consciousness is perceptual. Indeed, the very fact that emotions are conscious suggests that they belong to the class of perceptions, since consciousness seems to be restricted to perceptual states. The James-Lange approach to emotional consciousness is both intuitively plausible and theoretically advantageous. It allows us accommodate emotional consciousness within a unified account.

I think we can go one step further. It can be shown that at least one plausible theory of perceptual consciousness can be extended to apply to emotions. The theory I have in mind was first advanced by Ray Jackendoff (1987). Jackendoff begins with the observation that the senses are hierarchically organized. Low-level sensory systems respond to very discrete local features of an incoming stimulus, without integrating those features into a unified whole. In vision, low-level systems register edges, but not whole shapes. In hearing, the low-level captures discrete tones, but not melodies. Another stage of processing is required. Intermediate-level perceptual systems take these local features and integrate them: edges become contours, and tones become tunes. The intermediate-level delivers coherent representations, rather than a buzzing confusion, but it is not the final level



of perceptual processing. The intermediate level retains information about stimulus properties that is highly specific. If you see a chair, the intermediate level encodes its specific orientation with respect to your vantage-point. If you hear a word or a song, the intermediate level encodes information about its pitch and other acoustic properties. In order to achieve categorical recognition of a stimulus, it is often necessary to abstract away from such specific information. We need to be able to recognize a chair from many angles, and we need to be able to know what song is being played or what word is being uttered, not just the unique qualities of the performance or speech sounds we are hearing. High-level perceptual systems abstract away from specific features and facilitate recognition. A chair representation at the high level may be vantage-point invariant, and a high-level representation of a word may be invariant across the variable qualities of different speakers' voices. Given this organization, Jackendoff asks, where is consciousness? The answer should be obvious. When we see an object, we experience it from a point of view; when we hear a sound, we experience its unique acoustic profile. Consciousness resides at the intermediate level.

Can this story apply to emotions? Jackendoff (1987) implies that it cannot. After arguing that the account applies readily to the familiar senses, he makes no effort to argue that emotions are perceptual, much less that they are organized hierarchically. Instead, he characterizes emotions as 'markers' that colour the quality of visual, auditory, and other sensory experiences. This is puzzling for a reason mentioned earlier. Emotions obviously feel like something, and it would be odd to think that they get their qualitative character in a different way than other mental states that feel like something. It would be more parsimonious if we could explain emotional consciousness in the same way that we explain perceptual consciousness.

Elsewhere I have argued that Jackendoff's theory of perceptual consciousness can be extended to apply to emotions (Prinz 2004, 2005). Emotions, I have argued, are states in interoceptive systems, and interoceptive systems, like other perceptual systems, are probably hierarchically organized. There has been surprisingly little research on the organization of interoceptive systems, but there are principled reasons for postulating the following hierarchy. Interoceptive systems must register changes in each of the various bodily systems, including

circulatory, digestive, respiratory, muscular, and endocrinal changes. Each of these systems may have its own three-level hierarchy in the brain. In respiration, for example, the low-level may register each breath, the intermediate-level may register the specific quality of our breathing pattern, and the high-level may categorize these patterns (e.g., as strained or deep or relaxed). In addition to these processing hierarchies for each bodily system, the brain presumably has some way of recognizing when different bodily systems are working in concert; the brain can recognize patterns of activity across different bodily systems. This, too, may involve a hierarchy: first bodily changes in individual bodily systems are registered; then patterns across different systems are registered; and finally, the brain abstracts away from the specific details of a pattern to categorize it as a pattern of a specific type. These hierarchies are not separate. We can think of the pathway that looks for patterns as containing the pathways that monitor each bodily system. The pattern-detecting pathway monitors these bodily pathways and registers when different bodily states co-occur.

Now where do emotions fit in this picture? If emotions are interoceptive states that register patterns of bodily change, then emotions can be identified with the simultaneous and integrated occurrence of perceptions in pathways that register changes in specific bodily systems. Emotion categories (fear, sadness, anger, and so on) are probably applied when we abstract away from details about the specific way in which we are currently breathing or tensing our muscles. In each episode of fear, for example, our bodies will be configured somewhat differently, so to recognize fear as such, we may need to abstract away from those differences. Fear recognition occurs at a high level. But what about fear experience – the qualitative character of fear? Intuitively, fear experience is located at the intermediate level. When you experience fear, the quality is determined by specific features of your current bodily state. If you cower in a corner it will feel different than if you arch backwards. If you freeze, it will feel different from fleeing. If your breathing stops for a moment, that will feel different than if your merely suffer from shortness of breath. Each of these differences affects the character of a fear experience on any given occasion.

Thus, emotions seem to arise at the intermediate level of two hierarchies. They comprise intermediate-level representations in each of

the pathways that monitor bodily systems, and they occur at an intermediate level of a pathway that registers patterns of bodily change. One can think of emotional experiences as integrated, co-occurring, intermediate-level representations in pathways that register bodily changes. The experience of fear is an experience of a heart racing in a specific way along with a specific rate of breathing and a specific level of muscle tension, and so on. Emotions are experienced as bound episodes of specific changes throughout the body.

This proposal locates emotional consciousness at an intermediate level of perceptual processing. If I am right, then emotional consciousness is a case of perceptual consciousness, and perceptual consciousness can be explained in the same way across sense modalities. This lends support to the thesis that emotions are perceptions.

## V. Emotions Are Quasi-Modular

So far I have argued that emotions are states in our interceptive systems that represent abstract concerns and become conscious in just the way that paradigm instances of perceptual states do. Perhaps this is enough for concluding that emotions are perceptions. But opponents of that thesis might quibble. Emotions seem to interact with cognitive states in a way that makes them seem quite different from ordinary perceptions. We often have emotional responses as the result of a cognitive process. We think about a situation, and our cognitive assessments determine our emotional response. In paradigm cases of perception, there is no role for intermediating thoughts. We see the world and then judge, rather than the other way around. This apparent contrast between emotions and perceptions ties into the idea of stimulus dependence, which I introduced in my definition of quasi-modularity. To approach this objection, I will argue that emotions are quasi-modular. I will begin with the first four features of quasi-modularity and work up to the fifth condition, which is stimulus dependence. I will argue that the alleged contrast between perception and emotion is not sufficiently great to deny the thesis that emotion is a form of perception.

The first condition on quasi-modularity is functional specialization. Quasi-modules serve specific functions. It should be perfectly obvious that this is true in the case of emotions. Emotions are designed to pro-

vide information about our relation to the world. They represent concerns, and they do so in a way that captures our attention and influences our behaviour. By drawing attention to concerns, emotions disrupt our ongoing plans and force us to adopt coping strategies. Fear, for example, draws our attention to a danger, and thereby forces us to decide whether to engage in avoidance behaviour. There is, in other words, a characteristic functional role for the emotions. Emotions are triggered by inputs that bear on well-being and they send outputs to centres involved in planning and action-guidance.

The second attribute of quasi-modules is that they break down in characteristic ways. Compare emotion and vision. Damage to low-level visual areas can cause blindness because visual information cannot get into the system if these structures are compromised. Likewise, damage to the centres that allow bodily information to get into the central nervous system leads to a reduction in emotional experience (Chwalisz et al. 1988). In both cases, however, there can be a residual capacity for experience through mental imagery. When the primary visual cortex is destroyed, visual images can occur, and when the spinal cord is damaged, emotions can be experienced through bodily imagery.

Damage to intermediate-level visual areas leads to blindness as well because the intermediate level of vision is the locus on conscious experience. Similarly, emotional experience can be diminished or eliminated by damaging portions of the anterior cingulate cortex, in a condition called *akinetetic mutism* (Damasio and Van Hoesen 1983). The anterior cingulate probably isn't a low-level interoceptive area; it receives inputs from other medial cortical areas like the insula. This suggests that the anterior cingulate is the locus (or a locus) of intermediate-level emotion processing.

Finally, damage to centres associated with high-level vision do not eliminate visual experience, but they prevent a person from interpreting their visual states – a syndrome called *associative agnosia*. Likewise, there is an emotional condition called *alexithymia* in which a person has emotions but fails to identify them correctly. Such individuals may misidentify their emotions, or mistake emotional responses for mere bodily aches and twinges (somatization). Poor emotion recognition has been associated with reduced or abnormal function in rostral areas of the anterior cingulate and structures in the medial frontal

cortex (Lane et al. 1997; Berthoz et al. 2002). These structures may contain the correlates of high-level emotion processing. In sum, emotions can malfunction in characteristic ways, and these breakdowns mirror the kinds of disorders found in vision and other sense modalities.

The third characteristic of quasi-modules is automaticity. I don't think that perceptual systems are entirely automatic or always automatic. When we form visual images, or engage in visual search, we exercise a considerable amount of control over our visual systems. But, the visual system *can* function automatically, as when a visual state is triggered bottom-up without effort or intervention. Visual pop-out effects are a clear case of this. Likewise, emotions can be triggered automatically. If someone were to suddenly pull the chair from under you, you would experience fear, and that fear would require no effort or control. Likewise, seeing someone cry can induce sadness, and getting cut off by another car on the highway can trigger rage.

The fourth characteristic of quasi-modules is innateness. Quasi-modules can certainly contain many learned rules and representations. In vision, we store memories of previously perceived objects and scenes, and we develop visual skills through practice (e.g., chicken-sexing or interpreting ambiguous images). But vision capitalizes on a stock of innate rules and representations. There are species-typical, biologically prepared mechanisms for detecting colours and shapes. Visual contrast and constancy effects, motion perception, and depth perception all trade on innate rules and representations. Emotions are equally dependent on innate rules and representations. We have innate mechanisms for perceiving changes in the body, and, as mentioned earlier, specific bodily patterns are innately disposed to be triggered by specific stimuli. For example, looming objects and loud sudden noises trigger the bodily response that we experience as fear.

Turn now to the final characteristic of quasi-modularity: stimulus-dependence. At first, it might seem a trivial matter to show that emotions are stimulus-dependent. If the body is perturbed in a particular way, we will register that perturbation. The stimulus drives the response. But, granting this sense in which emotions are stimulus-driven, there is another sense in which emotions seem quite unlike paradigm cases of perception. Under ordinary circumstances, emotions are not induced by directly altering the body. Rather, there

is a prior mental event – a perception or a judgment pertaining to one's current circumstances. For example, fear might be elicited by seeing a venomous spider or by judging that there is a prowler in the house. These mental events trigger a perturbation in the body, which then generates the emotion. Qua response to body change, emotions seem to be stimulus-driven, but qua prior perception or judgment, their induction seems suspiciously endogenous. Emotions seem to require one step too many to qualify as perceptual. When we see something, there is no other sensory state between the stimulus that we see and the resultant visual perception. In the case of emotions, there is an intermediary. A danger out there must be perceived by some sense other than interoception before it can trigger the bodily response that causes the emotion. In this respect, ordinary cases of emotion elicitation are indirect. Does this show that they are not stimulus-dependent?

I think not. Even if we grant that the link between an emotion and its eliciting condition is indirect, we must admit that, under many circumstances, the link is nevertheless driven bottom-up. We see the spider; it causes a visual experience, which then causes the emotion. The emotion is not chosen. It is triggered by the stimulus. In these cases, at least, emotional response has a passivity that is highly characteristic of perceptual responses. Emotions are often just as stimulus-dependent as paradigm cases of perception.

Emotions also exhibit what I called perceptual trumping in section I. Consider a situation in which an emotion has been triggered bottom-up by a prior perception or judgment. Once the emotion has been triggered, it conveys relational information. Fear represents danger, for example. We can, after an emotion has been elicited, form a belief that conflicts with the content of the emotion. After fear is elicited, we can form the belief that we are not in danger. When this happens, the emotion does not simply evaporate. Darwin tells an anecdote about being startled by a puff adder, which lunged at him from behind a glass cage. Darwin knew he was safe, but he experienced fear. While watching horror films, we experience the same phenomenon. Likewise, music can make us feel sad even when we know there has been no loss in our lives, and tickling can induce delight even though none of our goals have been satisfied. Emotions are recalcitrant. They cannot simply be overridden by contrary beliefs. In these cases, emotions are

like optical illusions: they persist even when we know that they are misrepresenting the actual situation. This confirms my contention that emotions are stimulus-dependent, and their dependence seems to be very much like the kind of dependence we see in paradigm cases of perception. So, I think the fact that emotions are indirectly linked to their eliciting stimuli should not distract us from the fact that they behave in ways that are just like quasi-modular perceptual states. Trumping is a symptom of stimulus-dependence. In sum, then, emotions have all the marks of quasi-modularity.

## VI. Conclusion

I began this discussion with four properties that characterize paradigm cases of perception. Perception takes place in modality-specific input systems; perceptions represent things; perceptions can be consciously experienced; and perceptions are quasi-modular. I suggested that anything exhibiting all four of these characteristics deserves to be called a case of perception. The bulk of this discussion has been dedicated to showing that emotions exhibit all four characteristics, and it is therefore appropriate to think of emotions as a form of perception. There is a sense in which emotions allow us to perceive matters of concern, and they do so by registering changes in the body. This conclusion has implications for theories of emotion, and it also has implications for other theoretical domains in which emotions have been implicated. For example, it is sometimes suggested that moral judgments are emotional in nature. If that is right, and if emotions are perceptions, then we can make literal sense of the phrase “moral perception.” Just as we can perceive danger and loss, we may be able to perceive right and wrong. I leave that possible implication for another occasion.

## References

- Berthoz, S., E. Artiges, P.-F. Van de Moortele, J.-B. Poline, S. Rouquette, S. M. Consoli, and J.L. Martinot. 2002. Effect of impaired recognition and expression of emotions on frontocingulate cortices: An fMRI study of men with alexithymia. *American Journal of Psychiatry* 159: 961–67.

- Cannon, W. B. 1927. The James-Lange theory of emotion: A critical examination and an alternative theory. *American Journal of Psychology* 39: 106–24.
- Chwalisz, K., E. Diener, and D. Gallagher. 1988. Autonomic arousal feedback and emotional experience: Evidence from the spinal cord injured. *Journal of Personality and Social Psychology* 54: 820–28.
- Critchley, H. D., S. Wiens, P. Rotshtein, A. Öhman, and R. J. Dolan. 2004. Neural systems supporting interoceptive awareness. *Nature Neuroscience* 7: 189–95.
- Damasio, A. R. 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. New York: Gossett/Putnam.
- Damasio, A. R. 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt Brace.
- Damasio, A. R., and G. W. Van Hoesen. 1983. Emotional disturbances associated with focal lesions of the limbic frontal lobe. In *Neuropsychology of Human Emotion*, ed. K. M. Heilman and P. Satz. (85–110) New York: Guilford Press.
- Dretske, F. 1988. *Explaining Behavior*. Cambridge (Mass.): MIT Press.
- Fodor, J. 1983. *The Modularity of Mind*. Cambridge (Mass.): MIT Press.
- Jackendoff, R. 1987. *Consciousness and the Computational Mind*. Cambridge (Mass.): MIT Press.
- James, W. 1884. What is an emotion? *Mind* 9: 188–205.
- Kosslyn, S. M. 1994. *Image and Brain: The Resolution of the Imagery Debate*. Cambridge (Mass.): MIT Press.
- Lane, R. D., G. L. Ahern, G. E. Schwartz, and A. W. Kaszniak. 1997. Is alexithymia the emotional equivalent of blindsight? *Biological Psychiatry* 42: 834–44.
- Lange, C. G. 1885. *Om Sindsbevaegelser: Et Psyko-fysiologisk Studie*. Kjøbenhavn: Jacob Lunds. Reprinted in *The Emotions*, ed. C. G. Lange and W. James, trans. I. A. Haupt. Baltimore: Williams and Wilkins, 1922.
- LeDoux, J. E. 1996. *The Emotional Brain*. New York: Simon and Schuster.
- Marañón, G. 1924. Contribution à l'étude de l'action émotive de l'adrénaline. *Revue française d'endocrinologie* 2: 301–25.
- McGurk, H., and J. MacDonald. 1976. Hearing lips and seeing voices. *Nature* 264: 746–48.
- Phan, K. L., T. D. Wager, S. F. Taylor, and I. Liberzon. 2004. Functional neuroimaging studies of human emotions. *CNS Spectrums* 9: 258–67.
- Philippot, P., C. Chapelle, and S. Blairy. 2002. Respiratory feedback in the generation of emotion. *Cognition and Emotion* 16: 605–627.



- Prinz, J. J. 2000. The duality of content. *Philosophical Studies* 100: 1–34.
- Prinz, J. J. 2004. *Gut Reactions: A Perceptual Theory of Emotion*. New York: Oxford University Press.
- Prinz, J. J. 2005. Are emotions feelings? *Journal of Consciousness Studies* 12(8–10): 9–25.
- Prinz, J. J. 2006. Is the mind really modular? In *Contemporary Debates in Cognitive Science*, ed. R. Stainton. (22–36) Oxford: Blackwell.
- Prinz, J. J. 2007. All consciousness is perceptual. In *Contemporary Debates in Philosophy of Mind*, ed. B. McLaughlin and J. Cohen. (335–357) Oxford: Blackwell.
- Stepper, S., and F. Strack. 1993. Proprioceptive determinants of affective and nonaffective feelings. *Journal of Personality and Social Psychology* 64: 211–20.
- Strack, F., L. L. Martin, and S. Stepper. 1988. Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology* 54: 768–77.
- Wiens, S. 2005. Interoception in emotional experience. *Current Opinion in Neurology* 18: 442–47.
- Wolpert, D. M., and J. R. Flanagan. 2001. Motor prediction. *Current Biology* 11: R729–R732.