## *Chapter Two:*

# *Classical cognitive science and emotion*

> *Think how much stronger the self will be when it deliberately uses reason and judgement to form a decision. For the mind freed from passions is like a fortress, and there is nothing more secure in which to retreat and find unceasing sanctuary.*
>
> Marcus Aurelius, *The Meditations*

It may seem presumptuous that a discipline barely forty years old already prides itself on having a 'classical' form and various 'non-classical' variants. Yet this is exactly how cognitive scientists describe the theoretical diversity that currently characterises their field of study. In this chapter I outline the main features of classical cognitive science, and then discuss the classical approach to emotion.

### *2.1. Classical cognitive science*

By 'classical' cognitive science, I intend to refer to the general style in which the cognitive research program was pursued in its first decades, from about 1950 to 1980. This style was marked by a number of assumptions that have since been challenged by various sections of the cognitive science community. These assumptions included:

*(i)* *Domain generality*
*(ii)* *Internalism*
*(iii)* *Discreteness*

In the rest of this section, I will briefly describe what is meant by each of these terms. First, however, I want to make a few general points.

During the years 1950-1980, nobody spoke of 'classical' cognitive science. The assumptions of domain generality, internalism and digitality were so widely accepted by cognitive scientists that they seemed essential components of cognitive science itself. Only later, when various sections of the cognitive science community began to challenge these assumptions, did it become clear that that cognitive science was not, in fact, committed to them. At that point, it became useful to distinguish different *species* of cognitive science. Because the assumptions of domain generality, internalism and digitality had prevailed in the early days of cognitive science, it became common to refer to this set of views as marking the 'classical' form of the discipline. Those who challenged one or more of these assumptions could then refer their own approaches as 'non-classical'.

The classical assumptions of domain-generality, internalism and digitality were challenged, respectively, by those adopting evolutionary, situated, and analogue approaches to the study of the mind. These approaches therefore, are all 'non-classical' in one way or another. However, just because an approach is non-classical in one way does not mean that it has to be non-classical in every other way. Just because the evolutionary approach rejects domain-generality, for example, does not mean that it has to reject internalism and digitality. However, in the final chapter I argue that while the various non-classical approaches are *logically* independent of one another, there are good *theoretical* reasons of a non-logical kind that make the non-classical approaches natural allies. While they do not logically entail one another, the non-classical approaches can be combined to make up a single coherent species of cognitive science that we might refer to as 'integrated non-classical cognitive science'.

This said, it is now time to look at what the assumptions of classical cognitive science actually say.

*(i) Domain generality*

As I noted in the previous chapter, the principle explanatory strategy of cognitive science is functional decomposition. On this view, understanding the mind is like understanding the body in the sense that both involve 'carving nature at its joints'. Cognitive science is supposed to proceed by producing an 'anatomy' (or a 'map') of the human mind, whose components are discovered by taking a componential approach to designing artificial minds. If the artificial minds thus designed behave just like human ones, this is good grounds for assuming that the human mind is composed of similar functional units.

In adopting this view, cognitive science assumes that the mind can be partitioned into distinct parts. This view, which is sometimes known as 'faculty psychology' is an ancient view, going back at least as far as Plato.[1] There are various ways in which one could go about anatomising the mind, but cognitive scientists in the period 1950-1980 assumed that the best way of doing so was to divide the mind, first, into three parts: a perception system, a general reasoning system, and a motor-control

[1] Faculty psychology was challenged in the eighteenth century by Hume and other associationists, who argued that the mind is a homogeneous entity governed by a few general principles. If the associationists are right, then the viability of functional decomposition as an explanatory strategy in understanding the mind would be seriously in doubt

system. This was not a new idea either; it figured in many previous models of the mind, including one proposed by Freud.

In the nineteenth century, however, Franz Joseph Gall had proposed a rather different way of doing mental anatomy. According to Gall, the mind is composed of a large number of subsystems, each of which operates solely in a particular *domain*. Gall did not spell out in any detail the criteria for individuating domains, but one gets a rough intuitive idea of what he meant when one reads the list of subsystems he identified; there were, for example, subsystems for musical ability, moral reasoning, and the appreciation of beauty (see Fodor, 1983).

Gall's taxonomy of mental structures is clearly orthogonal to the taxonomy assumed by most early cognitive scientists. In that taxonomy, a single perceptual system took in all sensory stimuli, processed them, and passed them to a single reasoning system, which then decided, on the basis of all this information, what instructions to pass on to the motor control system. There is no specific sensory system dedicated purely to the analysis of music or beauty. The sensory system, like the reasoning system and the motor-control system, is *domain-general*.

In Gall's view, things are quite different. Each subsystem works relatively autonomously with input of a certain class. The music subsystem only processes musical stimuli, for example. Each subsystem has its own relatively independent means for perceiving, reasoning and initiating motion. Another way of putting this is that perception, reason and motor-control are *domain-specific*.

The choice between the two taxonomies is not a black and white one. In 1983, for example, Jerry Fodor proposed a hybrid taxonomy that combined elements of both. In *The Modularity of Mind*, he argued that the sensory systems and motor-control systems were domain-specific, but that the central reasoning system was domain-general (Fodor, 1983). Fodor's criteria for individuating domains, however, were rather different from those envisaged by Gall. In fact, they were just the same as those normally used to individuate the five senses (with the modification that language comprehension was regarded as a distinct perceptual system in its own right, a kind of 'sixth sense'). According to Fodor, the distinct sensory systems worked relatively autonomously, but then passed all their information to a single, domain-general central system where all the data were integrated. Fodor's mind is still, then, fundamentally domain-general in its centre and bulk. Despite the small gesture towards domain-specificity, Fodor retains the basic assumption of domain-generality that characterises the classical approach.

Fodor's suggestion opened the way for other cognitive scientists to argue for a view much more like Gall's. In particular, a number of evolutionary psychologists began to argue, in the late 1980's, that the human mind was composed entirely of domain-specific systems, each of which had evolved to solve a particular adaptive problem faced by our ancestors. This has since become known as the 'massive modularity hypothesis'.[2] I discuss this view in chapter three.

### *(ii) Internalism*

CTM states that minds are computers, but (as we saw in chapter one) nothing is a computer except with regard to some other, external system. Whenever we are confronted with a claim that something has a mind, therefore, we can always ask where, exactly, the boundary lies between the internal and the external systems. Where, in other words, are we to locate the input-output boundary?

In the case of humans and other brainy creatures, most cognitive scientists in the period 1950-1980 tended to locate this boundary at the junction between the central nervous system and the rest of the body. The mind, in other words, was thought to *supervene* entirely on the brain. The only *bona fide* psychological states, therefore, are those that can be individuated without reference to the rest of the world outside the brain. This view is sometimes known as internalism or individualism, though, like most 'isms', these words cover a multitude of sins and mean different things to different people. *The MIT Encyclopedia of the Cognitive Sciences*, for example, takes a slightly different view of internalism, defining it as the view that 'psychology in particular and the cognitive sciences more generally are to be concerned with natural kinds whose instances end at the boundary of the individual' (Wilson, 1999: 397). Equating the input-output boundary with that of the *individual* is clearly somewhat different to equating it with the *brain-body* boundary.

Instead of trying to resolve this ambiguity by futile discussions about the where exactly internalists *should* locate the input-output boundary, we can simply identify the common assumption underlying all the different formulations of internalism. Although they may disagree about *where* the input-output boundary is located, all internalists assume that this boundary is to be *identified* with some physical feature of the organism. It is *this* assumption, therefore, that should be taken as the essence of internalism.

[2] The term is due to Sperber (1994).

In the 1980s, a loose federation of cognitive scientists began to reject internalism in favour of a more 'situated' approach to the mind. They argued that the input-output boundary was a moveable feast. On their view, the question of where the input-output boundary is to be located will depend very much on the particular context of enquiry. The boundary of the *mind* is not to be simply *identified* with *any* physical boundary, whether that between the brain and the rest of the body or that between the body and the rest of the world. Some mental processes may well supervene entirely on the brain, or even on just one part of the brain. Others supervene on the brain plus part of the body. Others supervene on the brain, body, and parts of the world. In Andy Clark's arresting metaphor, the mind often *leaks* out of the brain into the body and the rest of the world (Clark, 1997). I discuss this view in chapter four.

*(iii) Discreteness*

The machines designed by cognitive scientists in the period 1950-1980 were almost all *discrete-state* machines. In such machines, the transitions from one state to another are like sudden jumps or clicks. The various states are sufficiently distinct and definite for the possibility of confusion between them to be ignored. There are no intermediate positions between one state and another.

Many of the first generation of cognitive scientists assumed that this was the *only* way to design intelligent machines. In 1976 Allen Newell and Herbert Simon consolidated this impression by arguing that all cognitive agents would turn out to be digital machines. This is the thrust of their 'physical symbol system' hypothesis (Newell and Simon, 1976: 85). Strictly speaking, this claim has nothing to do with discreteness. A digital machine is not *necessarily* a discrete-state machine. A digital machine is simply one that can reidentify things it has made positively and reliably (Haugeland, 1996: 9).

Now, it *might* turn out to be the case that only discrete-state machines can succeed in being digital, but, if so, this would be an empirical discovery about discrete-state machines – we cannot make any such inference on purely conceptual grounds. The concepts of being discrete and being digital are logically independent. Nevertheless, for some reason the two concepts have often been confused, with the result that Newell and Simon's physical symbol system hypothesis seemed to confirm the view that cognitive science was exclusively concerned with discrete-state machines.

In the 1980s, this view was challenged by a growing band of cognitive scientists who were interested in using *continuous* machines to model the mind. They used components with continuously variable rates of activation, and linked them together in networks resembling groups of interconnected neurons. The connectionist movement was not the only section of the cognitive science community to call for analogue models of the mind. In the 1990s, they were joined by others of a more theoretical bent who proposed that cognitive science could benefit from using the tools of dynamical systems theory.

Dynamical systems are not all continuous. There are discrete dynamical systems as well. But most proponents of dynamical approaches to cognition have tended to concentrate on continuous systems. It is therefore plausible to regard the connectionist movement and the dynamical approach to cognitive science as fighting on a similar front. Both call into question the assumption that cognitive science is exclusively concerned with discrete-state machines. For this reason, I discuss them both together in chapter five, which is concerned with continuous approaches to cognitive science.

### *Other features of classical cognitive science*

A number of other assumptions were shared by many of the first cognitive scientists apart from the three listed above. Many, for example, assumed that mental representations are stored in the brain in a rich language-like code (dubbed 'Mentalese' by Fodor) which is independent of any natural language like English or Japanese (Fodor, 1975). This is known as the 'language of thought' hypothesis', and is a particularly strong form of representationalism. It is possible to adopt weaker forms of representationalism, according to which thoughts are representations, but do not take a language-like, propositional form. Cognitive scientists working outside the classical tradition are generally more likely to espouse such a weaker form of representationalism, and a few even claim to reject the idea that thoughts are representations altogether (although it is not clear whether, in fact, they are really just objecting to the strong Fodorian version). One can then, with hindsight, see a commitment to the language of thought hypothesis as another of the distinguishing features of classical cognitive science.

I take it, however, that the language of thought hypothesis is less central to the classical approach than the three assumptions described above. Whole movements in cognitive science have arisen based on the desire to challenge these assumptions. I think, then, that there are good grounds for considering these assumptions to be the main diagnostic features of the classical approach. In the next section, I show how they inform the classical approach to emotion.

## 2.2. The classical approach to emotion

Despite the general reluctance to address the emotions among most of the pioneers of cognitive science, there were, even in the early days, a few lonely voices calling for a more inclusive research program that would embrace emotional factors as well as classical cognitive processes. In 1963, for example, Robert Abelson proposed that cognitive psychologists should move away from their focus on 'cold' logical processes and address 'hot cognitions' (Abelson, 1963). In 1967 Herbert Simon himself argued cognitive models should include emotions (Simon, 1967).

Such calls for a more inclusive research program forced cognitive scientists to face up to the dilemma posed by Hume's distinction between reason and the passions. Since nobody was prepared to argue that emotions were not *bona fide* mental processes, cognitive scientists were forced to reject Hume's thesis about the non-representational nature of emotion. Hence, those cognitive scientists interested in emotion attempted to reduce emotions to particular kinds of thought. The resulting research project came to be known as appraisal theory.

### *Appraisal theory*

Appraisal theory assumes that emotions are evaluations of current situations. The first proponent of this approach was Magda Arnold, whose pioneering book, *Emotion and Personality*, practically inaugurated the cognitive science of emotion (Arnold, 1960). The research program that grew out of Arnold's work attempted to discover the features of situations and events which cognitive agents use to arrive at an emotional evaluation. For example, some have suggested that a key aspect of the antecedent situation is whether it was self-caused or other-caused (Smith and Ellsworth, 1985). On the basis of this and other criteria, represented in terms of a simple decision-tree, a computer could analyse any situation (if presented appropriately) and decide which emotion it should respond with. That is, it could take a linguistic description of a situation as input, and generate the name of a particular emotion as output. Such a computer would clearly have explicit internal representations of emotions.

During the past decade, an increasing number of computer models of emotion have been designed along these lines. Most of them have been based on two appraisal-type theories of emotion: the Ortony Clore Collins model (henceforth OCC) and Ira Roseman's model (Ortony, Clore et al., 1988; Roseman, Antoniou et al., 1996). Both of these models were constructed with computers in mind, so both are relatively easy to implement in software. The process of designing such simulations can help to test the theories as well as stimulating new questions (Picard, 1997: 195).

Both the OCC model and Roseman's model categorise emotions on the basis of the cognitive appraisal that people make about eliciting conditions. In the OCC model, emotions arise from valenced reactions to situations consisting of the consequences of events (is it good or bad for me and for others?), the actions of agents (do I approve of my actions and those of others?), and the aspects of objects (do I like them or not?). A decision tree involving these questions leads to twenty-two different emotions, from joy and distress to gratitude and anger. Roseman's model uses slightly different parameters to generate a total of seventeen emotions.

Both of these systems are framed in terms of rules and so are relatively easy to implement in classical digital computers (symbol systems). The OCC model has never been implemented in full in any AI system, but simplified versions of it have been used to synthesise emotions in various applications. Tomoko Koda used the OCC model to simulate a restricted set of ten emotional facial expressions on poker-playing software agents (Koda, 1996, cited in Picard, 1997). Clark Elliot augmented the OCC model to twenty-six emotion types and used these as the basis of his 'Affective Reasoner' system (Elliot, 1994).

### *The classical approach to emotion*

Within the field of cognitive science, appraisal theory may rightly be called the 'classical' approach to emotion, both because it has dominated the cognitive psychology of emotion, and because it implicitly adopts all the main tenets of classical cognitive science. In line with classical cognitive science, appraisal theory takes an implicitly domain-general, internalist and discrete view of emotion.

#### *(i) Domain-generality and emotion*

In appraisal theory, there is no provision for distinct emotional subsystems. A single system analyses all the relevant features of the situation, and then computes a single emotion as output.

#### *(ii) Internalism and emotion*

Appraisal theorists assumed that emotional processes supervened entirely on the brain. The brain required sensory and proprioceptive inputs, of course, in order to generate an emotion, and then instructed the body to move in certain ways in accordance with the kind of emotion generated, but in their 'essence', emotional processes were neural, not physiological.

*(iii) Discreteness and emotion*

The decision-tree model of emotion employed by appraisal theory implies that emotions are discrete in a number of ways. Firstly, an emotion is either present or it is not; there are no intermediate levels of activation. Secondly, no attention is paid to the temporal features of emotion, such as how long it lasts, and how quickly it is triggered.

*Criticisms of the classical approach to emotion*

The classical cognitive approach to emotion rejected Hume's explication of the traditional distinction between reason and the passions, but failed to specify any other way of explicating this distinction. In so doing, the classical approach to emotion threatened to undermine the distinction altogether and thereby eliminate emotion from the taxonomy of mental processes. It was not long before critical voices were raised. One of the most prominent critics was the psychologist, Robert Zajonc, whose influential article 'Feeling and thinking: preferences need no inferences', argued that it was important to retain the ancient distinction between cognition and emotion (Zajonc, 1980).

Zajonc's arguments, however, were weakened by a serious equivocation; his use of the word *cognitive* was decidedly ambiguous. At some points in his article, the term is used to designate a set of mental processes that differ in important ways from other mental processes of an 'emotional' nature. However, the experimental evidence that Zajonc cites in support of this distinction in fact supports a rather *different* claim; namely, that the appraisal process preceding the experience of emotion is largely inaccessible to *conscious* introspection (Zajonc, 1980). Thus the word 'cognitive' is best construed, at other points in the article, as synonymous with the term 'conscious'. This is misleading, to say the least. Unless we simply wish to identify conscious/unconscious distinction with the cognition/emotion distinction (which seems a very unattractive option), we must acknowledge that Zajonc's experiment does not provide grounds for regarding cognition and emotion as distinct kinds of mental process.

Zajonc's article was based, in part, on a clever experiment that he conducted himself, which was an extension of earlier work he had done on the 'mere exposure effect'. This term refers to the fact that, when subjects are exposed to novel visual patterns, and then asked to choose whether they prefer these or similar patterns to which they have *not* been exposed, they prefer the pre-exposed ones. Mere exposure, in other words, is enough to create preferences. In the experiment described in the 1980 paper, Zajonc presented the visual patterns so quickly that subjects were unable to state accurately whether or not they had seen them before. All the same, the mere exposure effect was still there. Subjects gave all sorts of reasons for preferring the pre-exposed patterns, but nobody gave the correct reason (the pre-exposure).

This experiment dealt a severe blow to the method then used by appraisal theorists to investigate the emotions. This method consisted of asking people to introspect and figure out what had gone through their minds just before they had had some emotional experience. The appraisal theorists hoped thereby to find out the rules that relate antecedent situations to consequent emotions. Zajonc's experiment showed that this method was flawed because people were often wrong about the mental processes that caused them to have emotions.

Zajonc *could* have concluded that the appraisal processes preceding the experience of emotion are often inaccessible to conscious introspection. However, he went much further than this, claiming that emotional preferences could be formed without the aid of any cognitive processes at all. This is clearly a *non-sequitur*. As Joseph LeDoux points out, most processes that are considered prototypical examples of cognition also occur without any conscious awareness, so the absence of conscious recognition cannot be used to infer anything about the cognitive or non-cognitive status of emotional appraisal (LeDoux, 1998).

By pointing to the powerful role of unconscious factors in preferences, Zajonc provided a useful reminder that introspection is of limited use when doing research in cognitive psychology. Unfortunately, this important reminder was somewhat obscured by the fact that Zajonc perpetuated the old mistake of equating cognition with consciousness.

### *What Zajonc really meant*

As Paul Griffiths notes, the debate between Zajonc and his opponents about the relationship between emotion and cognition can sometimes seem like a mere semantic squabble, as if all it were about was the correct use of the term ‘cognitive’ (Griffiths, 1997: 25). However, there are more substantive issues at stake. No one denies that some kind of information processing must go on before an emotion can emerge in response to a given stimulus. The real question is whether or not this processing is of a kind sufficiently like the paradigm cases of unemotional thought-processes (such as deductive reasoning) to justify treating emotions as kinds of thought. Most emotion research in classical cognitive science has assumed a positive answer to this question. This is the assumption that Zajonc opposes.

As I noted in the previous section, mental processes are understood in classical cognitive science as bare computations. There is no provision here for distinguishing among different kinds of mental process, such as cognition and emotion. Zajonc is best construed as pointing out that the classical models of emotion do not provide a way of distinguishing emotion from cognition. If emotions are representational, what distinguishes them from thoughts, which are also representational? What Zajonc really meant was that we need to supplement the classical approach with new theoretical resources if we wish to do justice to emotion as a separate category of mental process.

At the time Zajonc wrote his paper, in 1980, the non-classical forms of cognitive science described in this thesis barely existed. In the years since then, the emotions have received much more attention from those working with non-classical approaches than they received from the pioneers of the classical approach. This suggests an intriguing possibility. Perhaps the non-classical approaches provide just the new conceptual resources needed for understanding the emotions that Zajonc was calling for in 1980. By providing new criteria for distinguishing between different *kinds* of computation, the non-classical approaches can help cognitive science find a new, non-Humean way of explicating the difference between cognition and emotion.

*The propositional attitude theory of emotion*

At the same time as cognitive scientists were developing appraisal theory, a group of analytic philosophers were working along similar lines. Like the appraisal theorists, these philosophers argued that emotions could be analysable purely as particular kinds of mental representation or thought. Specifically, it was proposed that emotions might simply be kinds of *judgement*. A concise statement of this view has been put forward by Richard Solomon:

> What is an emotion? An emotion is a judgement …[For example,] my embarrassment is my judgement to the effect that I am in an exceedingly awkward situation... an emotion is an evaluative (or a normative) judgement.
>
> (Solomon, 1977: 185, emphasis in original)

Now, most psychologists and philosophers accept that some emotions, at least, *imply* certain judgements. It is hard to imagine how someone could feel guilty, for example, unless they judge, perhaps unconsciously, that they have done wrong. Solomon's claim, however, is much stronger. He claims that an emotion *is no more than* its constituent judgements. Emotions, on this account, are entirely reducible to (kinds of) thought.

This approach to emotion became known, among analytic philosophers, as the 'cognitive' theory of emotion. This name is misleading, since it suggests a concern with the findings of cognitive psychology, when in fact there was, at that time, virtually no contact between the philosophical and psychological investigations into emotion. I will therefore follow Paul Griffiths in referring to this philosophical tradition as the 'propositional attitude theory of emotion' (Griffiths, 1997: 2).

Although there was no real dialogue between the cognitive psychologists who developed appraisal theory and the philosophers who pioneered the propositional attitude theory of emotion, the two theories are remarkably similar. Both reject Hume’s claim about the non-representational nature of emotion, and both attempt to reduce emotions to particular kinds of thought.

*Criticisms of the propositional attitude theory*

One of the most articulate philosophical critics of the propositional attitude theory of emotion is David Pugmire. In his book, *Rediscovering Emotion*, he argues persuasively that emotions cannot be reduced to judgements. His most powerful arguments turn on cases in which a person experiences an emotion of which the alleged constituent judgements are at odds with the person’s explicit beliefs. Pugmire dubs such emotions ‘irrational’, and provides a number of examples. One example concerns a man who does not believe in ghosts and yet is overtaken by fear as he enters a deserted house. Another is that of a bereaved widow who finds herself angry at her husband for having ‘left’ her, even though she knows full well that his death was an accident. Such emotions are not reducible to judgements, Pugmire claims, since the judgements they might be taken to consist of conflict with the person’s avowed beliefs.

Proponents of the propositional attitude theory of emotion have tried to deal with irrational emotions in a variety of ways. One is to say that the object of the emotion is represented less definitely in the person’s mind; the man entering the deserted house does not believe it is haunted, since he does not believe in ghosts, but he does think that there is a danger of some unspecified kind. Another is to say that the belief is not certain; the man does not believe firmly in ghosts, but nor is he certain that they do not exist, and he is better off safe than sorry. Pugmire argues that not all irrational emotions can be explained in these ways. Sometimes a person’s beliefs explicitly preclude the judgements alleged to be constitutive of the emotion, yet the emotion is still felt. The man explicitly states that he does not consider the house to be in the least dangerous, and yet he is still afraid when he enters it. So, Pugmire concludes, the propositional theory is in trouble.

Pugmire's objections to the propositional attitude theory fail, however, and for the same reasons as Zajonc's criticisms of appraisal theory: thoughts and judgements can be *unconscious*. Neither Zajonc's experiment with the mere-exposure effect, nor the cases of irrational emotion cited by Pugmire, rule out the possibility that emotions are reducible to *unconscious* thoughts. Like the appraisal theorist, the proponent of the propositional attitude theory can answer these objections by appealing to deeper cognitive resources offered by unconscious beliefs. When the man who is afraid of the deserted house tells us that he believes it is safe, there is no obvious reason why we should take his statement at face value. Perhaps he is simply *unaware* of his true beliefs. Or, perhaps he is only aware of *some* of his beliefs. If we take a domain-specific view of the mind, there could be various mutually inconsistent beliefs held by different mental parts, and it may be senseless to ask which of these beliefs is more truly *his*.

Pugmire is aware of the first of these responses, but not the second. When he points out that there are problems with tracing unconscious beliefs, then, he fails to see that this might be because we are assuming that the agent must have a single set of mutually consistent beliefs. Pugmire himself makes this dubious assumption in speaking of *irrational* emotions in the first place. He also assumes that, in cases of conflict between reported belief and a hypothetical unconscious belief, we must decide which of the two is the agent's *real* belief.

According to Pugmire, the claim that unconscious judgements must always have been at work in irrational emotions 'would rank as dogmatic, as an *ad hoc* hypothesis on the part of a psychoanalytically minded cognitivism' (Pugmire, 1998: 27). But this goes too far. Beliefs, or at least *representations*, are the fundamental explanatory tool of classical cognitive science. But cognitive scientists are not simply being dogmatic when they claim that that mental processes are thoroughly representational. There are good reasons that can be adduced to support this claim, some of which were outlined in the previous section. Among these reasons is the fact that we already have a good theoretical account of how representations can be processed by such obviously material entities as electronic computers, and thus a possible solution to the mind-body problem. In rejecting the representational theory of mind, Pugmire puts the mind-body problem back in the realm of mystery and thus leaves open the door to dualism.

Pugmire's main motivation for rejecting the representational theory of mind is his desire to save the Humean distinction between reason and the passions. He seems to think that, unless we can save this Humean thesis, we will have to give up the ancient distinction between cognition and emotion altogether. This, however, is clearly a *non-sequitur*. Just because we reject Hume's account of the distinction between reason and emotion does not mean that we have to give up the distinction altogether. There may be other ways of explicating this distinction, ways that do not appeal to a non-representational view of emotion. In the following chapters, I argue that the non-classical variants of cognitive science provide the resources that allow us to explicate the cognition-emotion distinction without giving up the representational theory of mind (and thus without giving up CTM).