

# Towards A Computational Theory of Mood

Laura Sizer

---

## ABSTRACT

Moods have global and profound effects on our thoughts, motivations and behavior. To understand human behavior and cognition fully, we must understand moods. In this paper I critically examine and reject the methodology of conventional ‘cognitive theories’ of affect. I lay the foundations of a new theory of moods that identifies them with processes of our cognitive functional architecture. Moods differ fundamentally from some of our other affective states and hence require distinct explanatory tools. The computational theory of mood I propose places them within the context of other mental phenomena and is consistent with the empirical data on moods.

- 1 *Introduction*
  - 2 *The affective spectrum*
  - 3 *Why the conventional approach cannot explain moods*
  - 4 *The computational theory of mind*
  - 5 *The foundations for a new theory of moods*
    - 5.1 *That moods are cognitively impenetrable*
    - 5.2 *That moods are operations of the functional architecture*
  - 6 *Conclusion*
- 

## 1 Introduction

Anyone who has ever experienced moods like melancholy, joy or anxiety can attest to the pervasive and sometimes dramatic effects they have on us. Our moods are sometimes no more than subtle accompaniments to our lives, merely tinting the edges of our thoughts. However, at other times they overwhelm us, consuming us with anxiety or ebullience or a melancholy we just can’t shake. The profound impact that moods have on our thoughts and behavior makes them an important subject of study for cognitive scientists and philosophers of mind and psychology.

Despite their significant effects on our mental lives, moods have largely been overlooked in the philosophical literature on affect. For the past thirty years ‘cognitive theories’ have dominated philosophical discussions of affect.

These theories, which I call the conventional approach, have focused almost exclusively on emotions, perhaps motivated by the assumption that moods and emotions are essentially similar. The unique properties of moods have proved troublesome for the conventional approach, however, and attempts to account for moods by stretching and modifying theories of emotion have failed.

These failures are instructive. They suggest that, contrary to conventional wisdom, moods differ fundamentally from paradigmatic emotions like instances of fear and anger. I argue that the conventional theories of affect cannot yield satisfactory explanations of mood and that a new and distinct explanatory methodology is required. Drawing on ideas and comments made by John Haugeland, Zenon Pylyshyn and Richard Davidson, I lay the foundations of a new, computational theory of moods. I argue that moods ought to be identified with modulations and biases in the operations of our cognitive functional architecture. My proposal liberates mood from its conventional association with emotion and provides a framework and direction for further theorizing.

## 2 The affective spectrum

In common language, use of affect terms like ‘emotion’ and ‘mood’ is inconsistent and fairly unprincipled. Unfortunately, the academic discourse does not fare much better. Indeed, the final analysis may conclude that these traditional categories should be discarded. The explanations for and mechanisms underlying different types of affective phenomena may not resemble the conventional folk psychological categories with which we are familiar.<sup>1</sup> None the less, I must begin the discussion here with some provisional description of the phenomena under investigation. It is best to present affective states as a spectrum of different types of experiences.<sup>2</sup> The spectrum spans from certain emotions (what I call paradigm emotion states) at one end of the spectrum to moods at the other. By moods, I mean generalized, nonspecific affective states like melancholy, ennui or ebullience. Different affective states are arranged along this spectrum based on two features: 1) the nature of and role played by the objects of the affective state, and 2) the degree to which the affective state is tied to beliefs and knowledge.

Certain instances of anger, fear, envy and jealousy exemplify *paradigm emotion states*. One reason I call these paradigm emotion states is that they are examples of emotions often used in philosophical discussions of emotion

<sup>1</sup> Griffiths ([1997]) reaches this conclusion regarding emotion. He argues that there is no identifiable psychological type picked out by the term ‘emotion’. Therefore the old category should be dissolved because it groups together very different, unrelated psychological states.

<sup>2</sup> My model is in some ways similar to and inspired by one presented by Armon-Jones ([1991]).

and seem to serve as models upon which many philosophical accounts of emotion have been based. As a specific example of a paradigm emotion state, imagine your experience of fear when you see that the stranger walking towards you late at night is holding a knife. You know that knife-wielding strangers are (or can be) dangerous and you feel you are in a threatened, perilous position. You have the strong desire to extricate yourself from this situation without getting hurt. As these thoughts race around in your head, your heart pounds, your palms sweat, and all of your attention is focused on the approaching stranger and the project of avoiding danger.

Paradigm emotion states are reactions to and are about particular, determinate objects. For instance, my fear when I encounter the armed stranger is about or directed toward the perceived stranger with a knife. The object of a paradigm emotion state is not necessarily a real, extant object; I could be hallucinating or mistaking the set of keys he is carrying for a knife. However, by saying that paradigm emotion states are object-directed, it is implied that there is something that the subject could pick out of his experience as the focus of attention, which is determinate and differentiable from the rest of his experience (Armon-Jones [1991]).

An important feature of paradigm emotion states is that they involve and are closely tied to our beliefs. Encountering a knife-toting stranger causes fear in most people, but encountering a teddy bear does not. This is because most of us have certain beliefs about knives—that they are, in the hands of strangers whose purposes we cannot divine, dangerous—that we don't have about teddy bears.<sup>3</sup> If someone were to react fearfully to a teddy bear, we would assume that he holds some unusual beliefs about teddy bears, or at least about this particular one (perhaps he believes there is a bomb in this one). If he did not hold some such belief, we would be quite right to question his sanity or wonder if he understands what it means to be afraid of something.

This fact reveals that there is a rational element in paradigm emotion reactions. By 'rational' I mean that the objects of paradigm emotion states and the beliefs we hold about them play the roles of reasons or justifications for our emotion reactions. It makes sense to ask why someone is afraid or angry, or whether their reaction is justified. One way of understanding this is that the objects of paradigm emotion states and the beliefs we hold about them often play reason-giving roles for the emotion reaction. This does not imply that emotion reactions are always rational, only that we look to the objects of our emotions as providing reasons for them. Another way to see the relationship between paradigm emotions and beliefs is that if the relevant

<sup>3</sup> It is not necessary that the cognitive components of paradigm emotion states be full-blown, fully conscious deliberations. They may simply be identifications or conceptualizations of the object that trigger our emotional reactions.

beliefs change, the emotions may change as well. If I realize that it is a set of keys and not a knife that the stranger carries, my fear will subside. Because of the different beliefs most of us hold about keys and knives, fear of a knife is usually warranted, but fear of keys (or a teddy bear) is not.

There are other types of affective experiences that we might call emotions, although they vary slightly from what I am calling paradigm emotion states. For example, we can fear more abstract objects like the Y2K bug, or the possibility of nuclear destruction. These affective states are quite similar to paradigm emotions in that they are closely tied to and affected by our beliefs. For example, if I learn and am convinced that the Y2K bug is a hoax manufactured by computer firms to generate business, then my fears will subside. Some affective states like irrational fears or phobias seem to be malfunctions or over-reactions of paradigm emotions. Like the paradigm emotions they are reactions to and are focused on determinate objects, events or states of affairs (snakes, clowns, enclosed spaces). However, they are more disconnected from and are relatively invulnerable to certain kinds of belief change and information (which is why they are termed 'irrational'). Irrational fears persist even if the person accepts and believes that there are no good reasons to be afraid. Joseph LeDoux ([1996]) argues that some fear reactions (particularly to evolutionarily salient stimuli, like the perception of a predator) occur without the mediation of higher cognitive processes, although higher cognitive processes can influence and censor the continuation of the response after the fact.<sup>4</sup> These 'automatic' emotion reactions are like paradigm emotions in that they are reactions to and are focused on determinate objects or states of affairs. They are connected to our beliefs to the extent that an automatic fear response to what is initially perceived as a snake will dissipate upon the realization that the object is a garden hose, not a snake. These sorts of affective states, and others, populate the middle of the spectrum. They share some but not all of the properties of paradigm emotions or moods, or possess them to a lesser degree.

At the far end of the spectrum are *moods*—nonspecific, indeterminate affective states such as melancholy, general anxiety, ennui and ebullience. These differ greatly from paradigm emotions in several respects. An example of a mood state is melancholy. Melancholy is usually marked by a pervasive, undifferentiated feeling of sadness, or a feeling of complete weariness and pessimism. As Haugeland ([1978], p. 223) describes it:

<sup>4</sup> There has been substantial debate in psychology over the relationship between emotion and cognition. For arguments that at least some emotion reactions do not require cognition see Zajonc ([1980], [1984a], [1984b]); Panksepp ([1994]); as well as LeDoux ([1996]). For defense of the view that emotions require cognition see Lazarus ([1982], [1994]), Lazarus, Coyne and Folkman ([1984]), Clore ([1994]), among others. Griffiths ([1997]) discusses this literature briefly and argues that automatic affective reactions comprise a distinct affective kind from affective states that involve higher cognitive processes.

The change from being cheerful to being melancholy is much more thorough and far-reaching than that from having a painless foot to having a foot that hurts. Not only does your foot seem different, but everything you encounter seems different. The whole world and everything in it, past, present, and future, becomes greyer, duller, less livable. Minor irritations and failings are more conspicuous and less remediable; ordinary things are no longer fun, lovely or pleasing [...] Likelihoods and improbabilities invert, as do what seems relevant to an issue and what seems beside the point [...] Many things affect our moods, but our moods also affect how things affect us; and in neither case is it quasilinguistic or rational. We do not state or believe our moods, or justify them on the basis of evidence or goals; they are just the way things are.

Moods do not have determinate objects and instead seem to permeate and affect all of our thoughts, desires, feelings and motivations. Moods are 'about' everything and nothing in particular. None the less, a subject in a mood like melancholy may complain about or dwell on a series of different issues, objects or events. Relationships, a job, taking the dog for a walk, losing the house keys, indeed anything that comes our way can take on the characteristics and colour of the mood. These objects do not cause our moods in the way that perceiving a knife-wielding stranger causes my fear. Instead the mood merely focuses on or takes up these objects or states of affairs.

As Haugeland points out, moods are typically disconnected from and relatively unaffected by our beliefs. The nonspecific feeling that nothing is going your way or that everything looks bleak is not tied to any specific beliefs about the world and is not easily affected by changes in our beliefs.<sup>5</sup> A melancholic or anxious mood will persist regardless or even in spite of what we believe about the world. Someone in an anxious mood may know that the pain in his temple is much more likely to be a headache than a brain tumor, but his worries about the possibility of a brain tumor will persist, even after the headache is gone.

My primary goal in defining this spectrum of types of affect is to present and describe the affective phenomena that I discuss throughout the paper. While it is important eventually to understand and explain the types of affective states that occupy the middle of the spectrum, this is not my current focus. We will be better equipped to deal with the middle-ground cases once

<sup>5</sup> A sudden switch from believing that everything looks bleak to believing that everything is rosy is a symptom of mood change, or mood swing, and not an example of a belief causing mood change. Coming to believe that what I thought was a knife is really a set of keys causes a change in my paradigm emotion state by replacing an object that I believe is dangerous with an innocuous one. However, if my mood changes, it can cause me to go from thinking that everything looks bleak to thinking that everything looks rosy. Here, the mood causes the belief change, and not the other way around.

we have a better understanding of the mechanisms responsible for those affective states that define the ends of the spectrum. I also wanted to illustrate that there are considerable differences between the types at each end of the spectrum. This is an important point to keep in mind in light of the fact that many philosophers who have written about affect have assumed that all affective states are essentially the same, and hence that one theory or explanatory apparatus can be applied to all types of affect.<sup>6</sup> In the next section I challenge this claim. I argue that conventional theories of affect (which primarily address states like the paradigm emotions) do not yield adequate explanations of moods.

### 3 Why the conventional approach cannot explain moods

Most of the philosophical work on affect has focused on states like the paradigm emotions, with the assumption that one theory or explanatory mechanism will apply equally well, *mutatis mutandis*, to all affective states. The conventional (cognitive) theories of affect tend to be tailored to the properties of emotions and assume that with a little tweaking they can be applied to explain moods. In this section I examine the cognitive or conventional theories of affect and argue that they cannot yield a satisfactory account of moods. The reasons the conventional approach fails are informative, as they not only support the conclusion that a different approach must be adopted to explain moods, but they indicate where we should locate such an effort. In what follows I discuss Anthony Kenny and Robert Solomon's applications of the cognitive theory to moods, as well as what can be called the 'neo-cognitive' theories of Claire Armon-Jones and Eric Lormand. I describe two difficulties that plague these accounts, which I refer to as the *circularity objection*, and the *global effects problem*. I argue that cognitivist attempts to account for the global effects and 'objectless' properties of moods inevitably lead to circular explanations.

For several decades cognitive theories of emotion have dominated philosophical discussion of affect. I refer to these theories as exemplifying the 'conventional' school of thought on affect because theirs has been the received view in philosophy and has shaped much of the current debate. The conventional, cognitive, theories are based on the idea that intentional, representational states are essential components of any emotion state. More specifically, they hold that emotions are essentially composed of propositional attitudes, which has led some to refer to them as 'propositional attitude theories' of affect. This view incorporates a folk psychological, representational theory of the mind, which sees the human mind as a representation

<sup>6</sup> Notable exceptions to this are Lormand ([1985]) and Griffiths ([1989], [1997]).

generating and manipulating device. Humans think or 'cognize' by representing information, beliefs, desires and goals, and then manipulating those representations according to certain rules of inference. Cognitive theories of affect are so called because they stipulate that (cognitive) intentional states such as beliefs, judgments and evaluations are the essential and identifying components of affective states.

Many versions of this approach (Kenny [1963]; Solomon [1976], among others) specify that epistemic propositional attitudes such as beliefs or judgments are essential components of affective states. Other theories in the cognitive tradition abandon the stipulation that epistemic propositional attitudes are necessary components and broaden their notion of cognition to include non-epistemic attitudes such as imagining and seeing-as (Stocker [1987]; Armon-Jones [1991]). Their approach is still 'cognitive', however, because it maintains the central claim that affective states are explained and identified by their constituent intentional, representational (cognitive) states.

Most cognitive theories acknowledge that physiological feeling elements play an important role in emotion experiences. Cognitivists like Lyons ([1980], [1992]) hold that the cognitive evaluations of stimuli cause the different physiological responses associated with certain affective states. However, since different emotions often have very similar physiological components, it is always the cognitive states that serve to individuate and identify the specific emotion.

The conventional approach seems to work well for paradigm emotion states like the instance of fear described in section two. My fear consists in the belief that there is a knife-wielding stranger approaching me, the belief that armed strangers are dangerous, and the evaluation of the situation as threatening. These intentional states, coupled with the desire to avoid danger, cause the physiological responses associated with fear: a pounding heart, sweaty palms, the motivation to run in the opposite direction. The emotion reaction is centered on the perception of the knife-wielding stranger and the evaluation or appraisal of the situation as one that is threatening or dangerous, and therefore undesirable. The experience of fear seems to require that something be evaluated or perceived as threatening or dangerous. Cognitive theorists draw upon this intuition and argue that to experience fear one must be thinking about something that it is threatening or dangerous.<sup>7</sup>

The conventional theories of affect have focused almost exclusively on emotions and therefore the majority of cognitive theories are, properly speaking, cognitive theories of emotion. This is because while most cognitive

<sup>7</sup> For cognitivists, one must at least be imagining or entertaining the possibility that something is threatening, dangerous. See Stocker ([1987]).

theories acknowledge the existence of affective states like moods, they do not think the differences between them and paradigm emotion states are very deep or significant. The conventional view of moods is that they are 'objectless' emotions. Emotions are object-directed intentional states; they are about or directed towards specific things. One is always in love with, angry at or afraid of someone or something. Moods, on the other hand, are 'objectless'. They are not essentially focused on specific things, events, or people in the way that paradigm emotion states are. One is not in a state of ennui, for instance, *about* anything.

The 'objectless' property of moods has presented a considerable challenge to the conventional view. Because the conventional approach is based on the assumption that affective states are constituted and defined by intentional mental states, it must contend that affective states, including mood states, have *some* sort of intentional object. Anthony Kenny's solution is to argue that moods have intentional objects such as 'everything' or 'things in general'. A similar approach is suggested by Robert Solomon, who argues that a mood like melancholy is about 'the world as a whole' or 'anything that comes our way':

We are often unaccountably depressed, on days when for no reason everything seems black; but pointless depression is not objectless depression, and the objects of depression are the things which seem black (Kenny [1963], pp. 60–1).

[...] moods are generalized emotions: an emotion focuses its attention on more-or-less particular objects and situations, whereas a mood enlarges its grasp to attend to the world as a whole, typically without focusing on any particular object or situation (Solomon [1976], p. 71).

Kenny and Solomon's solution does capture something about the global, pervasive effects of moods like melancholy and anxiety. These moods pervade all of our thoughts to an extent that we 'just feel depressed about everything', or judge that 'nothing is going our way'. However, their accounts fail to give a satisfactory means of identifying and explaining moods.

Kenny and Solomon's characterization of mood falls prey to the *circularity objection*. Upon closer examination of their accounts, we can see that the only thing doing any explanatory work is the stipulation that the subject is in a particular mood. In order to illustrate this point, let us say that Esther and Marvin are each in a different mood state, M1 and M2: Marvin is melancholic (M1) and Esther is ebullient (M2). When we apply Kenny and Solomon's solution to the problem of explaining 'objectless' affective states, we see that it defines Marvin's being in M1 in terms of his believing that everything looks black, or that he judges that everything is depressing. Esther, who is in M2, is enthusiastic about everything and believes everything



is a cause for excitement.<sup>8</sup> The problem with Kenny and Solomon's solution is that the only distinction between M1 and M2 is that Marvin is melancholic (about everything) and Esther is ebullient (about everything), which is exactly what we set out to define. Describing moods as having general objects like 'everything', 'things in general' or 'the world as a whole' does not yield explanations or means of identifying the states distinct from the stipulation that the subjects are in those states. Hence, their explanations are circular.

A possible response to the problem is to focus on defining the *attitude* melancholic Marvin takes toward 'the world as a whole'. We could therefore describe Marvin's affective state as an inclination to add 'It is sad that p' to all or most of the propositions, p, that Marvin holds.<sup>9</sup> Again, this description captures some of our experiences of a mood like depression but does not yield a satisfactory definition and identification of the state.

A satisfactory account of mood must explain the global, unbounded effects moods have on many, if not most, of our thoughts and attitudes. This is what I refer to as the *global effects problem*. Haugeland describes the problem for cognitivists in accounting for the global effects of mood in his article, 'The Nature and Plausibility of Cognitivism' ([1978]). He discusses the limitations of cognitivist explanations of mind and presents moods as an illustration of these limits, characterizing them as a 'potentially serious hurdle' to cognitivism. He points out that, '[i]f melancholy were an input representation ("melancholy here now") it would have to accompany and infect every other input, and transform the meanings of them all' ([1978], p. 223).

Because of the unbounded, global effects of mood, the 'p' in 'It is sad that p' can be filled in with all sorts of different propositions Marvin holds. This once again leads us into the circularity objection that Kenny and Solomon's solution encountered. In order to address the global effects problem, one must expand the set of mental states modified or affected as a result of Marvin's mood. But then the only thing doing the work of identifying Marvin's mood state is the disposition to add 'it is sad that p' for many or most p, which amounts to describing Marvin as sad, which is essentially what we set out to define.

This leads us to the next problem. There is much more to a mood like melancholy than thinking things are sad. We also feel weary, hopeless, doubt our abilities, believe negative outcomes are more probable than positive ones, and so on. What makes melancholy so pernicious is its effects on a wide range of our thoughts and attitudes. This presents a daunting challenge to the

<sup>8</sup> My point holds even if Esther and Marvin are ebullient/melancholic about 'most everything' or 'more things than not'. Any general object that is sufficient to capture the scope of a mood will run us into the same problem.

<sup>9</sup> Thanks to Elliott Sober for bringing this possible objection to my attention.

cognitivist: to capture the many aspects of a melancholic attitude without (circular) recourse to mood terms.

Cognitivists have not appreciated fully the difference in the relationships between paradigm emotions and their objects and moods and any 'objects' they may have. Paradigm emotions are responses to perceived objects or states of affairs that are to some extent causally responsible for the subject's experiencing the emotion. But with moods, it is the mood itself that pulls in 'whatever comes our way' as its 'object'. Moods saturate our experiences, turning the objects in our experiences into the objects of our moods. This is captured in Kenny's statement that 'the objects of depression are the things which seem black'. Moods do indeed imbue and affect *whatever* comes our way and therefore the grammatical or intentional objects of states associated with our moods are incidental to the identity or explanation of the mood. The problem is that, unlike emotions, there is nothing<sup>10</sup> that unites and explains the representational states associated with mood except the fact that the person is in that mood. Being in the mood is a prior fact to the having of certain representational states. But because there is nothing in the cognitive theory's explanatory arsenal that allows them to explain and characterize moods except intentional, representational states, cognitivist accounts of mood are inevitably circular.

Armon-Jones, who offers a 'neo-cognitive' theory of affect, analyzes the difficulties traditional cognitive theories encounter with moods and suggests a slightly different approach. She argues that an account of moods (which she refers to as 'M-states') must focus on 'patterns of thought' and 'modes of conceptualization' instead of on the semantic contents of object-directed mental states. Armon-Jones concludes as I have that the intentional contents of specific representational states seem incidental to the identity of a mood. She describes moods as 'undifferentiated affective frames of mind, identified by the imbuing of ongoing experience with a particular quality, rather than by the direction of experience upon an object' ([1992], p. 134).

Armon-Jones also describes what I have called the global effects problem. She argues that our moods range indiscriminately over items of experience that bear no logical connection to each other (except that they are 'whatever comes our way'), and that the rules of logical inference do not apply to the subject matter of general moods:

Though Tom construes his son, the car keys, the absent letter, etc., as irritating, these thoughts do not have a role as premises since they do not connect in such a way that some conclusion, say, in the form of an implied course of action, could be inferred from their conjunction ([1991], p. 101).

<sup>10</sup> By 'nothing' here I mean no object or locus of attention which is determinate and differentiable from the rest of the subject's experience.

Therefore, Armon-Jones argues, moods must be identified in terms of *patterns* of thoughts. She points out that ‘the metaphors of colouring and tonal quality capture the sense in which M-states derive their identity primarily from mode or manner of thinking rather than from what is thought’ ([1991], p. 85). In order to explain and identify particular moods we must explain their characteristic patterns of thought. She proposes that the patterns of thought associated with particular moods are characterized by certain modes of thinking/feeling, or what she terms ‘construing’. As she defines it, to ‘construe’ *x* as *y* is to construct an interpretation (*y*) of *x*, or to structure one’s experience by certain concepts ([1991], p. 31). Construals can consist of thoughts ranging from basic apprehensions of an object to recollections, or ‘imagination-based, “as-if” interpretations’ ([1991], p. 33). Furthermore construals can consist of thoughts that are undifferentiated with respect to objects.

Armon-Jones defines M-states in terms of these patterns or modes of construing: a particular M-state, *m*, is composed of a set of representational states which are all characterized by the mode of thinking/feeling or ‘construing’ *m*-ly. A joyful mood, therefore, is composed of representational states which all have in common the characteristic of being thought or ‘construed’ joyfully:

M-thought, though undifferentiated, projects a structured, unifying interpretation onto ongoing experience such that it coheres with the M-state [...] The construal ‘irritating’ projects an interpretation that structures the irascible agent’s undifferentiated experience of ongoing events [...] Like emotion thought, M-thought qualifies as cognitive because, as a mode of construing, it involves conceptualization of experience ([1991], p. 90).

Armon-Jones’ suggestion that an explanation of mood must focus on patterns of thought is a valuable one. Her notion of ‘construing’ is broad enough to encompass the many different kinds of thought associated with moods. However, it is not clear whether Armon-Jones’ solution escapes the problems inherent in other cognitive accounts of moods. To identify particular moods she must explain their characteristic patterns of thought. However, she is committed to the assumptions of the cognitive approach, and ultimately must explain moods in terms of sets of cognitive, representational states. Her only explanation of what unites a set of representations and identifies them as ‘joy’ is that the representations are products of a certain joyful mode of construing. Defining joy in terms of construing joyfully is patently circular and hence doesn’t really tell us anything about joy that we didn’t already know. Therefore her account does not avoid the circularity objection.

None the less, Armon-Jones' view contains a valuable insight in that it illuminates the source of the problem cognitive theories have with mood: what characterizes mood is *how* we think, not *what* we think. It suggests that what identifies and explains mood is not any set of intentional, representational states, but the processes that produce, or affect the production of those states. If this is the case, then an analysis of mood in terms of representational states (as the conventional view requires) misses the level at which the regularities of mood are most clearly exhibited and can be explained. In order to characterize and explain the patterns of thought associated with particular moods we need to examine the processes that are responsible for those patterns.

An analogy may help to clarify my point. Imagine that the pictures you take with your camera all come out fuzzy and blurry; once you develop them, the images of Aunt Ida, the dog at the beach, springtime in Madrid, etc. all look fuzzy. Now, to get to the bottom of what is happening with these pictures—to understand *why* your camera is producing blurry, fuzzy pictures—you don't look at the photographs themselves. Although the quality of the photos might indicate that *something* is amiss, the photos themselves won't tell you what it is—they won't explain the phenomenon. *Anything* the camera is aimed at will be represented fuzzily in the photograph. In order to identify and explain what is happening you must look at the underlying mechanisms that produce the photographs. In this case, it is a property of the underlying mechanisms (being out of focus) that explains the blurry quality of all the images.

Conventional philosophical theories of affect try to characterize moods in terms of their constituent representational states and fail. The reason for this failure is that they are attempting the equivalent of analyzing hundreds of blurry photographs in order to explain the quality they all share. The problem is that blurry photographs of Aunt Ida, the dog, and Madrid in spring have nothing in common except that they were all generated by a camera that was out of focus. Similarly, anxious or joyful thoughts about work, the bath mat, the neighbor have nothing in common except that the processes which generated them were biased in a certain way.

Eric Lormand's ([1985]) theory of mood takes an important step towards addressing this insight. Although Lormand's view is in some respects similar to Armon-Jones' and he places himself within the cognitivist camp, his theory is distinct from those discussed so far because he argues that moods are *not* intentional mental states. Lormand holds that we have more intentional states than are active (involved in deliberation, reasoning) at any one time. Moods are non-intentional variables ('sieves') that determine which of our intentional states are active at a given time. For example:

An explanation for why one's dog might seem to be more of a nuisance than usual when one is in a bad mood may be that one's mood causes one's beliefs about the unpleasant qualities of the dog to become more active than one's beliefs about the pleasant qualities of the dog ([1985], p. 399).

Lormand's theory makes a significant move away from the traditional cognitive theories in arguing that moods are non-intentional. He goes on to suggest a more complete explanation of what moods are, but presents it in terms of their relationships with intentional states. The resulting account as sketched resembles Armon-Jones' in that its characterizations of mood ultimately rest on and appeal to the sets of intentional states associated with (made active by) particular moods. Lormand states, for example, that 'an elated mood facilitates the activation of that set of evaluative beliefs which evaluate things as wonderful, and hampers the activation of those which evaluate things negatively' ([1985], p. 400). In order to avoid any charges of circularity Lormand must define the relevant belief sets without relying implicitly upon notions of the kinds of beliefs characteristic of different moods. Without this step Lormand's analysis of a mood like elation explains it as that which activates the set of beliefs that evaluate things as wonderful. If 'evaluating things as wonderful' is understood as a catch-all for the beliefs associated with elation, then the account is problematically circular. By tying the identity of moods to the intentional states they bring about, Lormand's view falls prey to the same problems that plagued other cognitivist theories of mood. A satisfactory account of mood must be able to characterize the processes identified with different moods independently from the intentional states they bring about.<sup>11</sup>

In this section of the paper I have argued that the conventional, cognitive theories of affect are unable to adequately explain mood. I have deliberately remained agnostic about the cognitive theory's ability to explain successfully other affective states such as the paradigm emotions, because the focus of this paper is mood, not emotion. However, my conclusion may be interpreted as helping proponents of the cognitive approach to the extent that it 'gets them off the hook' for explaining moods. Much of the criticism of the conventional approach has focused on its inability to account for the 'objectless' affective states. Proponents of the cognitive approach may use my arguments to claim that they don't have to explain moods on the grounds that moods are the product of altogether different cognitive mechanisms.<sup>12</sup>

<sup>11</sup> See Griffiths ([1989]) for a different discussion of Lormand's theory.

<sup>12</sup> I do not think that this is enough to save cognitive theories of emotion. The research suggesting that some of our emotion reactions proceed 'under the cognitive radar' for instance (see fn. 4), is a serious challenge to the cognitive theory. See Griffiths ([1997]) for an overview of the problems with cognitive theories of emotion.

I have suggested that conventional explanations of mood fail because they do not target the level of cognitive organization at which the operations governing moods are manifested and at which the explanatorily useful regularities of mood are exhibited. In the remainder of the paper I elaborate this claim and propose a computational theory of mood. The computational theory of mind, which situates representational, intentional processes in the context of other functionally defined cognitive systems and processes, gives us some of the explanatory tools needed to account for moods. The wide-ranging effects of moods on our representational states and processes suggest that the source of these effects is change at a deeper level of cognitive organization. The theory of mood I present accounts for the global effects of mood without running into the circularity problem.

#### **4 The computational theory of mind and Pylyshyn's cognitively penetrable/impenetrable distinction**

Before I propose my theory of moods I must briefly review the ideas and concepts central to my argument. These will be familiar to those versed in computational theories of mind. Computational theories of mind conceive of human cognition as an interaction of many complex systems operating at different levels of organization or abstraction. Different levels of cognitive organization exhibit certain regularities not evident or explainable at other levels. I argue that the regularities associated with moods are exhibited at a deeper level of cognitive organization than the one on which the conventional views focus. Moods are best understood as the results of processes at the level of our functional architecture. This explains their global effects on our representation level states and processes while avoiding the circularity objection. The circularity objection is avoided by the fact that one can identify and describe the processes associated with moods independently of their effects on representation level states.

The computational theory of mind draws on the idea that humans act on the basis of internal, mental representations: beliefs that the Eiffel Tower is in Paris, desires for a decent croissant, fear of spiders, and so on. However, it situates representational processes in the context of other functional and physically specified levels of organization. The computational view holds that these representations are instantiated physically in the brain as cognitive symbols or codes. Operations carried out on these symbols nomologically correlate with changes in and the semantic processing of representational content, and thereby cause us to act in certain ways.

The mind is composed of many complex systems in operation at different levels of organization. Physical (neural, biological) structures and principles interact to form functionally specifiable structures. These make up the

functional or syntactic level of organization including, among other things, the functional architecture. The *functional architecture* defines the basic operations, resources and constraints of the system (such as rules for storing, retrieving and manipulating information and other symbols, and the constraints imposed by limitations on memory and attention resources) as well as the cognitive control structure. The control structure directs and assigns resources and dictates which processing rules are to operate at a given moment. The functionally specifiable structures instantiate and allow for the intentional, representational states and processes we typically think of as 'cognition'. At this semantic, intentional level of organization are processes and mental states which are explainable in terms of the semantic content of representational states (like beliefs, desires, and goals) and the rational, inferential rules governing the interpretation and manipulation of those states.

The systems functioning at each of these different levels of organization require different methods of explanation or levels of description. While each level is autonomous from the others, the levels are all interrelated. For instance, we can describe a process at the representational level without discussing its physical realization. However, the two levels are related in that the representational level process is instantiated or realized in the hardware. The hardware, therefore, imposes certain constraints on representation level functions. A complete understanding of cognition requires an understanding of how all the levels work together.

Pylyshyn ([1980], [1984]) has introduced a means of distinguishing between representational and functional level processes. He describes the distinction between representational and functional level processes as 'fundamental', in that

it marks the boundary in theory between two highly different kinds of principles—the kind stated in terms of semantics, in terms of properties of the *intentional objects* or *contents* of thoughts or other representations—and those that can be explained in more traditional, functional terms ([1984], p. 262; italics in original).

One of the hallmarks of representation-governed processes is that they are *cognitively penetrable*. A function or process is cognitively penetrable if its performance depends on and is explained in virtue of what the system represents (the beliefs, goals, and knowledge of the agent). In some ways, this is simply further explication of the notion of a representation-governed process. However, it also offers a method for deciding if a particular cognitive phenomenon is a representational or functional level process. Processes of the functional architecture are cognitively *impenetrable*; we can explain them

without *any* reference to semantic rules and representations and solely in virtue of the fixed functional capacities of the system.

Cognitively penetrable processes and functions display a certain amount of *plasticity*; changes in beliefs and goals affect representational level processes in rationally explicable, or meaningful ways. Cognitive penetrability and plasticity are characteristics of representation-governed cognitive processes because they demonstrate that the performance of the process or generation of the state is sensitive to the semantic contents of representations. Therefore unlike physiochemical or reflex reactions, or certain processing and inference rules, certain aspects of human behavior are determined by and sensitive to changes in represented information. We can explain these sorts of regularities only by reference to the content of those representations.

Pylyshyn is not saying, however, that every aspect of a representational level process is completely plastic. Relatively immutable rules of inference and decision-making necessarily operate over and guide the manipulation of our representations. Pylyshyn holds that there are certain principles like rationality that guide and constrain transitions between representational states in terms of their content. The notion of rationality allows us to draw principled connections between beliefs, goals, and intended actions and is required to explain and make coherent much of our behavior. Processes attributable to operations of the functional architecture are those that are explicable without *any* reference to representations, tacit knowledge or inference.

The functional architecture and representation level processes do interact in principled, defined ways. Pylyshyn argues that they interact through ‘well-behaved’ interfaces: the specific inputs and outputs of *transducers*. A transducer translates or maps physical events in one form to another in accordance with consistent rules for doing so. Transducers are part of the functional architecture because they perform purely physical processes in virtue of their intrinsic physical, chemical, biological properties. A transducer function is such that whenever it receives a certain kind of input it produces a particular output, regardless of the context; the behavior of a transducer is not altered by cognitive processes.

In order to be part of the fixed functional architecture, a process or function must be cognitively impenetrable. However, Pylyshyn is careful to specify that the capacities of the functional architecture are only ‘relatively’ fixed. The functional architecture is not absolutely immutable; changes in biology, chemical, and hormonal systems may affect its operations. By ‘fixed’ Pylyshyn means fixed with respect to certain types of influences. The functional architecture is affected by influences *of the right type*: those that are not ‘under a cognitive description’:



When I say that some effect is produced by an input ‘under a cognitive description,’ that is to be read as an abbreviation for the statement that the relevant generalization is exhibited (and hence some particular observed regularity is explained) only when the input is so described [ . . . ] [W]hen input-output functions can be systematically altered by ‘purely cognitive means,’ which is to say by antecedent events under a cognitive description, then we have prima facie evidence that the function in question is governed by epistemic representations (i.e. goals and beliefs), and hence is not part of the functional architecture. ([1980], p. 155)

## **5 The foundations for a new theory of moods**

Now it is time to apply the concepts I have just discussed to the task of explaining moods. I concluded section three by arguing that the conventional view, which limits itself to representational level principles and explanations, cannot adequately explain moods. I will argue that moods are best explained as biases in operations at the level of our functional architecture. To make my point I will first demonstrate that the properties that differentiate moods from paradigm emotions also conform to Pylyshyn’s descriptions of cognitively impenetrable processes. While representational states are important components of an explanation of paradigm emotion states, moods require explanation in terms of the vocabulary and principles at the functional level of cognitive organization. I go on to demonstrate that viewing moods as part of the functional architecture explains many of the properties of moods that cause trouble for the conventional approach.

### **5.1 That moods are cognitively impenetrable**

Paradigm emotion states are examples of mental phenomena that exist at the representational level of cognitive organization. Explanations of paradigm emotion states must appeal to the semantic contents of representational states because paradigm emotions draw upon and depend on what the system represents. With the paradigm emotions, how we react emotionally depends on how we evaluate something or what beliefs we hold about it. For example, my anger towards a taxi driver who cuts me off in traffic is explained in part by my beliefs that such behavior is reckless, rude and dangerous. It is because I represent the driver’s actions as rude and reckless that I become angry, beeping my horn and yelling. By appealing to the semantic contents of representational states as well as principles like rationality we are able to make coherent the regularities of emotions and emotion produced behavior. Paradigm emotion states, therefore, bear the hallmarks of what Pylyshyn calls cognitive penetrability.

Moods, on the other hand, exhibit the features of cognitively *impenetrable* mental phenomena.<sup>13</sup> As I discussed in sections two and three, one of the unique characteristics of moods is that, unlike paradigm emotions, they are not focused on determinate, particular objects or states of affairs. They take up general objects like ‘everything’ or ‘anything that comes our way’, or focus momentarily on a series of unrelated objects or states of affairs (this is another way of saying that the focus of the mood is anything that comes our way). As I argued previously, these features of moods suggest that what we think about in particular plays no central identifying or explicative role in mood; our moods can focus on everything, nothing, or a whole string of unrelated items. Moods, therefore, operate independently from what the cognitive system represents.

Moods are disengaged or disconnected from our beliefs and knowledge, demonstrating that they are not penetrable or influenced by the semantic contents of our representations. If one is depressed about everything, or things in general, then no particular piece of new information or change in belief is going to affect this underlying pervasive mood. Even if one is depressed or anxious about many different things—the noise outside the window, the pain in one’s temple, or the upcoming meeting—the underlying state of anxiety remains regardless of the content of the subject’s worrying. New information (that the noise outside was caused by the cat) does nothing to relieve the anxiety; it simply shifts it to a different focus. These properties of mood conform to Pylyshyn’s description of cognitively impenetrable mental phenomena.

Some might argue that our thoughts, beliefs, and goals clearly influence our moods and therefore moods *are* cognitively penetrable. Certainly learning of a friend’s death or the frustration of a life-long goal has some influence on our moods. However, to establish that moods are cognitively penetrable, one must not only show that beliefs can influence our moods, but that the influences are *of the right kind*.

What this means is that, *first* one must demonstrate that the cognitive *input* is influencing mood in virtue of its semantic content. Pylyshyn states that, ‘[w]hen I speak of penetration as being cognitive, I mean that the effect produced is explainable in terms of the content of the penetrating belief or goal alone’ ([1980], p. 157). *Second*, one must be able to explain the nature of the influence (*output*) in terms of the content of the penetrating input. Showing a correlation between a new belief and a mood change is not sufficient to prove that the cognitive influence is of the right kind—that the *semantic content* of the specific belief or goal is doing the work.

<sup>13</sup> Pylyshyn anticipates some of what I say in his closing comments of *Computation and Cognition*, but does not develop the idea.

It is easy to imagine that some beliefs could have an effect on mood, especially catastrophic news like that of the death of a friend. The counter-argument being considered, then, is this:

- 1) If the semantic content of a belief can directly affect a change in mood, then moods are cognitively penetrable.
- 2) Acquisition of the belief that a friend has suddenly died can directly affect changes in mood.
- 3) Therefore, moods *are* cognitively penetrable.

What is at stake here is my claim that moods are cognitively impenetrable processes attributable to the functional architecture. If moods are cognitively penetrable processes, then they are not operations of the fixed functional capacities of the mind.

In response to this challenge I deny premise two above, and maintain that the influence that even catastrophic news has on mood is *not* of the right kind. The influence of the cognitive input can be explained in terms of a mediating transducer stage, which provides a principled non-cognitive interface between events under a cognitive description and any changes to operations of the functional architecture. Why ought one to accept my explanation over the counter-argument's claim that the cognitive state directly influences mood? I shall give two reasons, corresponding to the two steps stated above.

One reason is that there is unlikely to be a regular, systematic relation between the semantic content of the input and the consequent change in mood. There are principled, predictable relations between what we believe and our paradigm emotions. However the connection between our moods and our intentional, cognitive states is tenuous at best. As I discussed previously, moods characteristically are unaffected by and disconnected from our beliefs. This suggests that the semantic content of the belief alone cannot explain the mood change, and that changes at a deeper level of cognitive organization are responsible for changes in mood.

Some researchers suggest that while emotions are responses to particular objects or states of affairs, moods are responses to the overall state of the subject. Several researchers (Nowlis and Nowlis [1956]; Jacobson [1957]; Thayer [1989]; Morris [1989], [1992]) have argued that moods are part of a non-cognitive self-regulatory or self-monitoring system. We have limited resources of physical and mental energy with which to meet the demands of an ever-changing environment. We need not only to track the environment and the changes it presents, but also to monitor our own resources in relation to external demands. Morris, for instance, argues that moods play a role in this system by moderating our processing functions in relation to the resources available.

If this is the case, then moods are responses to the overall state of the subject and not to the semantic contents of particular beliefs. This explains why individuals vary in their mood reactions in the wake of catastrophic news. Beliefs may affect moods, but indirectly by influencing the overall state of the subject via a mediating transducer stage. A transducer stage is necessary in order to 'translate' information about environmental pressures and physical resources into a format ready for processing by a self-monitoring system. Upon learning about something tragic, like the death of a friend, people will experience emotions like grief and mourn the loss, but not everyone will experience a global change in mood.

A more compelling reason to favor my explanation over the counter-argument's is that the semantic content of the purportedly penetrating belief is not sufficient to explain the global nature of mood change. In order to establish that mood is cognitively penetrable, one must explain the effects of mood in terms of the semantic contents of the penetrating cognitive input. Therefore one must demonstrate that, at the other 'end' of the process, the nature of the influence or outcome is explained by the semantic contents of the penetrating cognitive input.

Moods affect a wide range of our thoughts, feelings and attitudes in ways that are not constrained by subject matter or inferential rules. This is the 'global effects problem' I discussed in section three. Haugeland's point about the difficulty of explaining the global effects of mood applies here as well. In order to establish that moods are cognitively penetrable phenomena one must explain how the semantic contents of a particular cognitive state could 'infect and change the meaning of' every other cognitive state in order to bring about the diffuse effects associated with mood. This does not seem possible at the semantic, intentional level, since it is not clear what sort of semantic level processing rule would allow for this. Widespread, global changes or patterns at the representational level are more likely to be effects of general changes in processing strategies than responses to the semantic contents of particular beliefs.

Learning of the sudden death of a dear friend will certainly have a profound impact on anyone. One might experience emotions like grief, anger, or regret. In addition, there may be a change in mood. The question, however, is whether the change in mood follows directly from the newly acquired information about the friend's death, or whether any mood change is the result of a complex collection of cognitive and non-cognitive factors. I conclude that cognition affects mood, but only indirectly (via transduction). Therefore, in Pylyshyn's terms, moods are cognitively impenetrable processes. This account allows us to acknowledge that cognitive states do have some (indirect) influence on our moods, while explaining the global and pervasive effects of mood on all our subsequent thoughts and feelings.

In order to prove that this is the right explanation, research will have to show us which biological or neurochemical properties are doing the work at the transducer stage. Our present understanding of the biological and neurochemical operations associated with mood is not sufficient to resolve this question for once and for all. While it therefore remains possible that certain beliefs can have an influence of the right kind on mood, given the factors I discuss above, this seems unlikely.

## **5.2 That moods are operations of the functional architecture**

Recognition that moods are cognitively impenetrable is a first step towards seeing them as processes at the level of functional architecture. I must now show that the regularities associated with moods are best understood in terms of processes at the functional level of cognitive organization. I made some preliminary steps toward establishing this in section three where I argued that the representation level vocabulary and principles of conventional theories of affect cannot capture the regularities and patterns of thought associated with mood. In order to explain the patterns or biases in thoughts associated with mood we must examine the processes that govern the creation and manipulation of representational states. Explaining moods in terms of changes in functional level operations allows us to do this. Furthermore, empirical data on mood suggest that differences between positive and negative moods are the results of biases in the operations of basic processes such as memory, categorization, and attention allocation. These data lend further support for my view.

Understanding moods as operations at the level of functional architecture allows us to account for their global effects without running into the circularity objection. The functional architecture consists of rules and heuristics for storing and retrieving information from memory, allocating attention and other resources, and includes the control structure which dictates the processing rules in effect at a given time. Changes in operations at this level have global effects on representation level states and processes but in a way that is independent from the semantic contents of those states. The changes are bottom-up and therefore influence the creation and performance of all affected representations and representational processes. This is what explains the generalized, content independent effects of moods.

Davidson ([1994]) suggests that, '[t]he primary function of moods [...] is to modulate or bias cognition. Mood serves as a primary mechanism for altering information-processing priorities and for shifting modes of information processing' ([1994], p. 52). Such functions and mechanisms are usually ascribed to the functional architecture, particularly the control structure. Empirical research shows mood-related effects on such basic cognitive

processes as memory, attention allocation and concept formation and application. Because the observed mood effects are global and indiscriminate with respect to the specific contents of thought, these data suggest that moods operate by affecting the processing rules or heuristics governing these processes. Changes in the operations of these processes would cause global changes in our cognitions. Therefore we can (in principle) identify particular moods with specific modulations and biases in the functioning of these processes. It is possible to specify and describe these biases independently of their subsequent effects on representation level states. This will allow us to formulate non-circular explanations of mood.

The empirical data currently amassed on mood furnish us with a rough picture of the sorts of biases in processing associated with positive and negative moods. Many studies have demonstrated mood-related effects on memory, perception, categorization and attention.<sup>14</sup> For example, our moods affect what and how we remember (Clark and Teasdale [1982]; Mineka and Nugent [1995]). One phenomenon that is familiar to many of us is ‘mood congruent recall’; the recall of certain affectively toned material is facilitated by a match between the current mood of the subject and the affective tone of the recalled material. By affecting what and how we remember, moods can influence how we categorize and understand what we perceive (Isen [1984]; Morris [1989]).

Studies also show mood-related biases in what is attended to and how information is processed that are unrelated to information content (Sinclair and Mark [1992]). For example, subjects in positive moods tend to take in a wider range of information and categorize more broadly than those in negative moods; they do not focus on the smaller details, but instead perceive ‘the bigger picture’ (Isen [1984]; Isen and Daubman [1984]). When subjects are in positive moods they make associations between ideas and categorize stimuli faster. Isen *et al.* (Isen, Daubman and Nowicki [1987]; Isen [1993]) argue that when people are in positive moods their brains are more flexibly organized and they tend to make more creative and unusual associations between ideas. Negative moods, on the other hand, are associated with a more narrow focus and greater concentration on detail. Judgements and categorization proceed more cautiously and draw more heavily on the evidence at hand. This results in slower and more careful processing of information.

Understanding moods as processing biases at the level of functional architecture explains why moods are ‘objectless’ or are experienced as applying generally to anything and everything. Moods are not themselves representational states, but are changes at a deeper level of cognitive

<sup>14</sup> See Thayer ([1989]) or Morris ([1989]) for an overview and discussion of these studies.

organization. The results of changes at this deeper level are widespread, affecting all of our thoughts, attitudes and motivations, but can also be subtle, explaining the sometimes vague, nebulous character of moods. To paraphrase Haugeland's conclusion about cognitivist accounts of mood: moods permeate and affect all kinds of representational states and processes, and yet are not representational themselves ([1978], p.223).

In the closing chapter of *What Emotions Really Are*, Griffiths ([1997])<sup>15</sup> outlines a computational theory of mood that is structurally similar to the view I defend. He argues that moods are 'higher-order functional states which determine which of a range of lower-order functional descriptions a person occupies' ([1997], p. 253).<sup>16</sup> In this way moods 'cause global changes in propensities to occupy other states and to respond to stimuli' (*ibid.*).

Griffiths' model allows for moods to dispose a person to a wide and diverse range of thoughts, emotions and behavior, and is consistent with the idea that moods are associated with *patterns* of thought and behavior. However, Griffiths' computational account, couched in terms of nested orders of functional states, needs to be more specific about the distinctions and relationships between higher-order functional states (moods) and lower-order functional states. In particular Griffiths should clarify how his model allows for lower-order functional states like emotions to affect higher-order states like moods. The theory I present is consistent with Griffiths' model, but avoids the sometimes awkward discussion of higher and lower orders of functional states. The biases in operations of the functional architecture with which I identify moods correspond to Griffiths' higher-order functional states and do dispose the subject to realize other, lower-order functional states. However on my account it is possible to specify these biases independently of the lower order states to which they dispose us. This allows for a clearer and more informative description of the relationships between moods, emotions and other mental states.

Griffiths also suggests that moods can be identified with particular neurochemical states. A virtue of Griffiths' account is that his implementation level model of mood mirrors and complements his computational model. He argues that 'a mood state is a neurochemical condition which modifies the propensities of one neural event to bring about another' ([1997], p. 256), and thereby alters a person's overall functional description. Griffiths suggests that particular moods might be identifiable with particular neurochemical conditions. I agree that an explanation of the relevant neurochemistry is necessary for a complete understanding of moods, and we should view a computational model as being realized by neurochemistry (at least in

<sup>15</sup> See also Griffiths ([1989]).

<sup>16</sup> Griffiths explains that a functional state is higher order in the sense that, given certain conditions, it disposes the person to realize one of various other (lower-order) functional states.

humans). However, research efforts to establish relationships between particular neurochemical factors such as levels of biogenic amines, cortisol and thyroid hormones has so far failed to find strong correlations between these factors and different mood states. Some research suggests that moods are implemented by a heterogeneous collection of neurochemical types, or perhaps are associated with fluctuations or failures in any of a number of regulatory mechanisms.<sup>17</sup> This suggests that a model of mood at the computational level is better able to capture certain properties associated with our everyday conceptions of mood and explain the unique effects of different mood types like melancholy, anxiety and joy. None the less, continued research into the neurochemistry underlying mood will deepen our understanding of moods, and must be allowed to inform as well as realize a computational model.

There is still a great deal we do not know about moods. The data I have outlined are representative of the research accumulating on this topic. These data are consistent with my proposal that moods operate at the level of functional architecture and not at the level of representational states. The processes governing memory encoding, storage and retrieval, categorization, and attention allocation are all cognitively impenetrable processes that are part of the functional architecture. The operations of these processes may be altered in accordance with certain functional level rules or heuristics but are not influenced by our cognitive, representational states.

## 6 Conclusion

Applying Pylyshyn's distinctions, I have argued that while paradigm emotion states are cognitively penetrable, moods are cognitively impenetrable mental phenomena attributable to operations of the functional architecture. More specifically, I identify moods with biases and modulations in the operation of processes such as attention allocation, memory retrieval, and categorization. This view challenges the conventional assumption that all affective states are essentially similar and underscores the need for a new and distinct method of explaining and understanding moods. The theory of moods I propose explains their unique properties and is consistent with the empirical data. It not only provides a structure for explaining moods, but also clarifies why the conventional view goes wrong.

Throughout the paper I have focused on the differences between paradigm emotion states and moods. But there are undoubtedly strong connections and interactions between these states. It seems clear, for instance, that particular moods increase the probability of experiencing certain emotions, and that

<sup>17</sup> See the chapter by Paula Schnurr in Morris ([1989]) for a discussion of this literature.



some emotions, particularly strong ones, lead to changes in moods. What I have tried to do here is make a convincing case for the view that paradigm emotion states and mood states differ significantly and require distinct explanatory approaches. Recognizing these differences will put us in a better position to understand how paradigm emotions, moods and other affective states interact. My proposal holds clear implications for further research and theorizing about all types of affective states.

### Acknowledgements

I thank R. Andreasen, B. Bunk, N. Carroll, B. Enç, L. Shapiro, E. Sober, D. Stampe, C. Stephens and two anonymous referees for their invaluable comments on earlier versions of this paper.

*Department of Philosophy  
University of Wisconsin-Madison  
5168 H.C. White Hall, 600 N. Park St.  
Madison, WI 53706  
USA  
lpsizer@students.wisc.edu*

### References

- Armon-Jones, C. [1991]: *Varieties of Affect*, Toronto: University of Toronto Press.
- Armon-Jones, C. [1992]: 'Affect, Objects and Rationality', *Journal for the Theory of Social Behaviour*, **22**, 2, pp. 129–43.
- Clark, D. M. and Teasdale, J. D. [1982]: 'Diurnal Variation in Clinical Depression and Accessibility of Memories of Positive and Negative Experiences', *Journal of Abnormal Psychology*, **91**, pp. 87–95.
- Clore, G. [1994]: 'Why Emotions Require Cognition', in R. Davidson and P. Ekman (eds), *The Nature of Emotion*, New York: Oxford University Press, pp. 181–91.
- Davidson, R. [1994]: 'On Emotion, Mood and Related Affective Constructs', in R. Davidson and P. Ekman (eds), *The Nature of Emotion*, New York: Oxford University Press, pp. 51–55.
- Griffiths, P. [1989]: 'Folk, Functional and Neurochemical Aspects of Mood', *Philosophical Psychology*, **2**, 1, pp. 17–30.
- Griffiths, P. [1997]: *What Emotions Really Are*, Chicago: University of Chicago Press.
- Haugeland, J. [1978]: 'The Nature and Plausibility of Cognitivism', *Behavioral and Brain Sciences*, **2**, pp. 215–60.
- Isen, A. M. [1984]: 'Toward Understanding the Role of Affect in Cognition' in R. S. Wyer and T. K. Srull (eds), *Handbook of Social Cognition* 3, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 179–236.
- Isen, A. M. [1993]: 'Positive Affect and Decision Making', in M. Lewis and J. Haviland (eds), *Handbook of Emotions*, New York: Guilford Press, pp. 261–77.

- Isen, A. M. and Daubman, K. A. [1984]: 'The Influence of Affect on Categorization', *Journal of Personality and Social Psychology*, **47**, 6, pp. 1206–17.
- Isen, A. M., Daubman K. A. and Nowicki, G. P. [1987]: 'Positive Affect Facilitates Creative Problem Solving', *Journal of Personality and Social Psychology*, **52**, 6, pp. 1122–31.
- Jacobson, E. [1957]: 'Normal and Pathological Moods: Their Nature and Functions', *Psychoanalytic Study of the Child*, **12**, pp. 73–113.
- Kenny, A. [1963]: *Action, Emotion and Will*, New York: Humanities Press.
- Lazarus, R. [1982]: 'Thoughts on the Relations Between Emotion and Cognition', *American Psychologist*, **37**, 9, pp. 1019–24.
- Lazarus, R. [1994]: 'Appraisal: The Long and Short of It', in R. Davidson and P. Ekman (eds), *The Nature of Emotion*, New York: Oxford University Press, pp. 208–15.
- Lazarus, R., Coyne, J. and Folkman, S. [1984]: 'Cognition, Emotion and Motivation: The Doctoring of Humpty-Dumpty', in K. Scherer and P. Ekman (eds), *Approaches to Emotion*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 221–37.
- LeDoux, J. [1996]: *The Emotional Brain*, New York: Simon & Schuster.
- Lormand, E. [1985]: 'Toward a Theory of Moods', *Philosophical Studies*, **47**, pp. 385–407.
- Lyons, W. [1980]: *Emotion*, Cambridge: Cambridge University Press.
- Lyons, W. [1992]: 'An Introduction to the Philosophy of the Emotions', in K. T. Strongman (ed.), *International Review of Studies on Emotion 2*, New York: John Wiley & Sons Ltd., pp. 295–313.
- Mineka, S. and Nugent, K. [1995]: 'Mood-Congruent Memory Biases in Anxiety and Depression' in D. Schacter (ed.), *Memory Distortion: How Minds, Brain and Societies Reconstruct the Past*, Cambridge, MA: Harvard University Press, pp. 173–93.
- Morris, W. [1989]: *Mood: The Frame of Mind*, New York: Springer-Verlag.
- Morris, W. [1992]: 'A Functional Analysis of the Role of Mood in Affective Systems' in M. S. Clark (ed.), *Emotion: Review of Personality and Social Psychology 13*, Newbury Park, CA: Sage Publications, Inc., pp. 256–93.
- Nowlis, V. and Nowlis, H. [1956]: 'The Description and Analysis of Mood', *Annals New York Academy of Sciences*, **65**, pp. 345–55.
- Panksepp, J. [1994]: 'A Proper Distinction Between Affective and Cognitive Process Is Essential for Neuroscientific Progress', in R. Davidson and P. Ekman (eds), *The Nature of Emotion*, New York: Oxford University Press, pp. 224–26.
- Polyshyn, Z. [1980]: 'Computation and Cognition: Issues in the Foundations of Cognitive Science', *Behavioral and Brain Sciences*, **3**, pp. 111–69.
- Polyshyn, Z. [1984]: *Computation and Cognition: Toward a Foundation for Cognitive Science*, Cambridge, MA: MIT Press.
- Sinclair, R. C. and Mark, M. M. [1992]: 'The Influence of Mood State on Judgment and Action: Effects on Persuasion, Categorization, Social Justice, Person Perception, and Judgmental Accuracy', in L. Martin and A. Tesser (eds), *The Construction of Social Judgments*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 165–93.

- Solomon, R. C. [1976]: *The Passions*, Garden City: Anchor Press.
- Stocker, M. [1987]: 'Emotional Thoughts', *American Philosophical Quarterly*, **24**, 1, pp. 59–69.
- Thayer, R. [1989]: *The Biopsychology of Mood and Arousal*, New York: Oxford University Press.
- Zajonc, R. B. [1980]: 'Feeling and Thinking: Preferences Need No Inferences', *American Psychologist*, **35**, 2, pp. 151–75.
- Zajonc, R. B. [1984a]: 'The Interaction of Affect and Cognition', in K. Scherer and P. Ekman (eds), *Approaches to Emotion*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 239–45.
- Zajonc, R. B. [1984b]: 'On Primacy of Affect', in K. Scherer and P. Ekman (eds), *Approaches to Emotion*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 259–70.