Journal of Experimental & Theoretical Artificial Intelligence,
Vol. 17, No. 4, December 2005, 371–384

The implications of advances in research on motivation for cognitive models

ARTHUR B. MARKMAN*, W. TODD MADDOX and GRANT C. BALDWIN

Department of Psychology, University of Texas, Austin, TX, 78712 USA

(*Received 22 June 2005*)

There has been an upsurge of research in psychology on the interface between motivation and cognition. Much of this work has focused on elucidating the structure of the motivational system, although this work has also begun to examine the influence of motivation on preference, choice and learning. The growing body of data provides an opportunity for computationally minded researchers to extend existing cognitive models to incorporate insights about the nature of the motivational system. This paper reviews some recent research and draws out the implications of this work for computational cognitive science.

Keywords: Motivation; Regulatory focus; Cognitive modeling

1. Introduction

Cognitive science has typically been concerned with information processing. The computational view of mind that has driven much of the work in the field posits rich internal mental representations that support sophisticated perceptual and reasoning abilities. Consequently, many of the greatest triumphs of cognitive science have come from collaborative research projects that have proposed computational models of the way information is processed and have tested those models with studies of human information processing. For example, research on analogical reasoning has developed complex models of the structural alignment process and has tested those models empirically (see, for example, Gentner *et al*. 2001 for a review of many of these efforts).

Over the past ten years, cognitive scientists have begun to express their dissatisfaction with the progress made by the field, suggesting that perhaps the computational

*Corresponding author. Email: markman@psy.utexas.edu

Journal of Experimental & Theoretical Artificial Intelligence
ISSN 0952–813X print/ISSN 1362–3079 online © 2005 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/09528130500283915

view of mind needs to be overturned (e.g. Thelen and Smith 1994, Port and Van Gelder 1995, Markman and Dietrich 2000a). At the heart of many of the critiques of traditional cognitive science methodology is the recognition that the focus on information processing emphasizes high-level cognitive processes over perceptual and motor processes (Glenberg 1997, Barsalou 1999, Prinz 2002). The danger of ignoring perception and motor action is that the representations that govern these processes might also affect the structure and content of the mental representations used by higher-level cognitive processes (like analogical reasoning) that have often been the focus of research in cognitive science. For example, many studies in which people search for a visual target in space have been explained with concepts like attention and memory. There is good reason, however, to believe that people's performance is strongly guided by low-level perceptual mechanisms without involvement of higher-level cognitive processing (Geisler and Chou 1995).

In the spirit of those critiques that have suggested that cognitive science attend more to perception and motor action, we focus on motivation, which we take to be another under-represented area in the field. In this paper, we begin by giving a brief definition of key concepts relating to motivation. Then, we review some initial evidence that motivational states influence cognitive processing in humans. Finally, we discuss the implications of this work for computational cognitive science.

2. Key aspects of the motivational system

Agents act in the world. In order to act effectively, they must have an end-state to pursue as well as an impetus to achieve that end-state. In the psychology literature, the end-states are called *goals*, and the impetus to satisfy a particular goal is called *motivation*. Both are required in order for an organism to act.

Goals have been a part of many cognitive models since the dawn of the cognitive revolution. For example, Miller *et al.*'s (1960) classic work assumed that agents have goals that they wish to pursue. Early AI planning systems like STRIPS planned a series of actions by taking a desired end state and breaking it down into a series of subgoals that could then be achieved (Fikes and Nilsson 1971). Anderson's important ACT family of models contains explicit representations of desired end states that are used to select the next production rule to fire (Anderson 1983, 1993). In all of these models, however, goals are simply representations of states to be achieved. Often, traversing from an initial state to a goal state is done by searching through a space of possibilities to find a set of intermediate steps that creates a path from the initial state to the goal.

In this paper, we focus on recent research that explores the structure of the motivational system and the influence of motivation on learning. We focus on this work for three reasons. First, this research demonstrates how activation of the motivational system influences both behaviour and learning. Second, much of this work demonstrates that significant aspects of the motivational system are not consciously accessible, and so the contents of the motivational system are unlikely to be manipulated explicitly by the cognitive system. This work contrasts with older studies of the influence of task motivation on performance, where the rewards associated with succeeding or failing on a task are explicit. Third, this work helps to bridge the gap between cognitive processing and perception and motor action by suggesting ways that information processing can be used to drive action.

A central tenet of current theories of motivation is that people typically do not have direct conscious access to many aspects of their motivational system (Bargh and Chartrand 1999, Markman and Brendl 2005). A person's goals can be primed in a variety of ways leading them to perform actions without direct awareness of why they are performing those actions. For example, Aarts *et al*. (2004) found that people who read about or observed someone else performing an action unconsciously adopted that goal. In one study, men who read about someone flirting with a woman were more helpful to a female confederate later in the study than were people not so primed. The participants in this study were not aware that their actions were affected by reading a story about flirting.

In another study, Nisbett and Wilson (1977) showed people an array of four pairs of socks and were asked to select the best pair. Unbeknownst to the participant, all of the pairs of socks were the same. People had a strong tendency to select the right-most pair. Presumably, they selected this pair, because they scanned the array from left-to-right, and then selected the pair examined most recently. People were unaware of this basis for their choice. When asked why they made the selection they did, they justified their choice in terms of the quality of the socks (see also Kruglanski *et al*. 2002).

People's lack of conscious access to key aspects of their current motivational state has important implications for cognitive models of action. To the extent that goals are incorporated into cognitive science models at all, they tend to be included as explicit representational structures that the cognitive system can reason about. While people clearly have representations of desired end-states they would like to achieve, many aspects of the motivational system are not consciously accessible (Markman and Brendl 2005).

A second key aspect of the motivational system that has been the focus of recent research concerns types of goals. It has long been known, of course, that goals may refer either to desired outcomes to be achieved or undesired outcomes to be evaded (Miller 1959, Carver and Scheier 1998). Goals with positive end states are called *approach goals*, while goals with negative end states are called *avoidance goals*. Approach and avoidance goals are not only distinct conceptually, but also they likely involve quite different neural structures in humans and other animals (Fanselow 1995, Pickering and Gray 2001).

The approach and avoidance systems are separate, because different kinds of actions are appropriate for approaching desired states and avoiding undesired states. Indeed, highly noxious elements in the environment can engage the sympathetic nervous system leading to 'fight or flight' behaviours. There is also some evidence that activation of approach and avoidance goals follows a different schedule. As the psychological distance to a goal state narrows, the activation of avoidance goals increases more steeply than does the activation of approach goals (Miller, 1959, Markman and Brendl 2000). This structure is useful, because it leads the motivational system to focus on avoiding undesirable states when there are noxious elements in the environment. However, when undesirable states are psychologically distant, it permits exploration of potential positive states (Carver and Scheier 1998).

More recent work has expanded our understanding of approach and avoidance motivation by developing the concept of *regulatory focus* (Higgins 1997, 2000).

Regulatory focus is a state of motivational readiness that may influence the way an individual acts and processes information. There is a state of readiness for approach called a *promotion focus*, and a state of readiness for avoidance called a *prevention focus*. Individuals with a promotion focus are sensitive to potential gains in the environment. They seek positive outcomes. Successfully obtaining a positive outcome leads to emotions of happiness, and failing to obtain positive outcomes leads to emotions of sadness. In contrast, individuals with a prevention focus are sensitive to potential losses in the environment. They seek to avoid negative outcomes. Successfully avoiding a negative outcome leads to emotions of relief, while unsuccessfully avoiding negative outcomes leads to emotions of anxiety.

An important observation from this research is that people's performance tends to be best when there is a fit between an individual's regulatory focus and the current circumstance. For example, people experience the greatest fluency when evaluating items that are consistent with their regulatory fit (Higgins *et al*. 2003). For example, Higgins *et al*. (2003) found that people with a promotion focus valued consumer products more when they focused on what they would gain by owning the products than what they would give up by not owning them. In contrast, people with a prevention focus valued consumer products more when they focused on what they would give up by not owning them than what they would gain by owning them. Presumably, this difference in preference is driven by the fact that positive characteristics of a product are easier to process for people with a promotion focus than for people with a prevention focus, while the reverse is true for people with a prevention focus.

This section is not intended to present a complete model of the motivational system (see Carver and Scheier 1998, Pickering and Gray 2001, Kruglanski *et al*. 2002 and Markman and Brendl 2005 for more comprehensive discussions). Instead, we simply wished to introduce a few complexities about the nature of motivation that arise from the psychology literature whose implications for cognitive science need to be explored in more detail.

Much of the work described so far has been focused on elucidating the nature of the goal system. This work has focused on whether goal activation needs to be conscious, how goals are connected to each other, factors that affect the activation of goals, and the existence of promotion and prevention regulatory foci. Research has only begun to explore the implications of motivation on information processing. This link between motivation and cognition is an important area of future research, and we turn to it in the next section.

3. Exploring the effects of motivation on cognition

There are a number of ways that activation of the motivational system can influence cognitive processing. The most straightforward is for the activation of goals to increase the accessibility of goal-related concepts. These accessibility effects would facilitate goal satisfaction by priming goal-related concepts. There is some evidence for these kinds of semantic priming effects that we review in this section (Patalano and Seifert 1997, Fishbach *et al*. 2003, Shah and Kruglanski 2003).

In addition, the state of the motivational system can influence learning. For example, people often learn categories in the context of performing other tasks

(Markman and Ross 2003). These tasks may lead to potential gains or losses for the learner. Consequently, the learner's regulatory focus can influence responding to tasks involving gains and losses, which in turn can affect what is learned. We also address these effects below.

3.1. *Goals and knowledge accessibility*

In order to pursue a goal, an agent must be in a state of readiness to satisfy a goal. One aspect of this state of readiness is a cognitive preparation in which concepts related to the active goal are made temporarily more accessible in order to facilitate goal satisfaction. This accessibility can be accomplished by assuming that goals exist in a semantic network in which nodes representing end states of goals are connected to nodes representing subgoals and means for achieving goals (Anderson 1983, Kruglanski *et al*. 2002). On this view, the activation of a particular goal would permit activation to spread to related subgoals and to means for achieving a goal.

In addition to a spreading activation mechanism of this type, there must be at least two inhibitory mechanisms. First, when a goal is satisfied, it must be deactivated so that other goals can be pursued. Second, there are often competing goals. For example, a person might want to eat an ice cream sundae but also stay slim, or might want to purchase a car, but also have enough money to go to college. Under some circumstances, activating one goal may lead to the inhibition of potentially competing goals, which should influence the accessibility of information relating to those goals. In this section, we discuss evidence that relates goal activation to accessibility.

Goals are associated with means for carrying out those goals. When an individual has a goal, but has not yet satisfied that goal, means for achieving the goal can remind people of the active goal. In one set of studies, Patalano and Seifert (1997) had people adopt particular goals and were given a plan to satisfy the goals. When they later saw words describing objects found in a dorm room, objects relating to the active goals reminded people of those goals.

As another demonstration of this point, Shah and Kruglanski (2003) asked participants to solve anagrams presented on a computer screen. Prior to beginning this task, participants saw a word flashed to the screen. For half of the subjects, this word was the name of a strategy for solving the anagrams; for the other half, it was a word unrelated to solving anagrams. Compared to those given the unrelated prime, subjects primed with a strategy name were better at solving anagrams. In addition, all participants performed a lexical decision task in which they determined whether strings of letters presented on the screen were words. The words in the lexical decision task were either related to solving anagrams or were unrelated. Participants who were given the strategy name as a prime were also faster to respond to anagram-related words in the lexical decision task than were participants who were given a neutral prime. Taken together, these results suggest that means for satisfying a goal (e.g. the name of a strategy) can prime the goal (e.g. to solve anagrams), which in turn can beneficially influence performance.

Active goals can also increase the accessibility of items related to the goal. For example, Markman *et al*. (in preparation) studied habitual smokers who had either a high or a low need to smoke. A Stroop task was used to assess the accessibility of items that varied in their relatedness to smoking. In the Stroop task (1935), people are shown a word printed in a colour. The task is to identify the colour of the printing. When the concept named by the word is accessible, people are slower to name the colour than when it is not accessible. In Stroop's original studies (Stroop, 1935), the words were printed in colours, and the words themselves named colours. Stroop found that people were particularly slow to name the colour of the printing when the colour of the printing was different from the colour named by the word. Much recent work, however, has used the Stroop task as a measure of accessibility for concepts beyond colour words (e.g. Forster *et al.* 2005, Jacobs *et al.* 2003)

In our study, people with a high need to smoke were slower to identify the colour (relative to those with a low need to smoke) when the words named brands of cigarettes or items that can be used to help someone smoke (such as lighters or ashtrays). They were faster to identify the colour when the words named consumer products unrelated to smoking (such as a DVD player). A similar pattern was obtained for peoples' ratings of their preferences for the items. Subjects with a high need to smoke gave reliably higher preferences for brands of cigarettes and items instrumentally related to smoking and expressed lower preferences for smoking unrelated items than did subjects with a low need to smoke.

Goals and means activate each other most strongly when the individual has a goal that they wish to achieve, but have not yet achieved (Forster *et al.* 2005). To demonstrate this point, Forster *et al.* (2005) asked people to watch a computer screen for a picture of a pair of eyeglasses and to notify the experimenter when the picture of the glasses was immediately preceded by a picture of a pair of scissors. During the task, participants also had to perform a lexical decision task for words that were related and unrelated to eyeglasses. Prior to seeing the scissors/eyeglasses combination, people were significantly faster to respond to the words relating to glasses than to the unrelated words. In the block of trials just after seeing this combination, people were significantly slower to respond to the words relating to glasses than to the unrelated words. Finally, late in the experiment, people were about equally fast to respond to the items related and unrelated to glasses. This pattern of data suggests that the active goal to search for glasses increased the accessibility of concepts related to the active goal. Once the goal is fulfilled, it is actively inhibited. When the goal is inactive, it no longer influences concepts related to it.

The data described so far suggest that goals and means are connected so that active goals can increase the accessibility of means related to those goals, and means related to unfulfilled goals can increase the accessibility of those goals. The goal system also helps people avoid temptations that might hinder goal satisfaction. In particular, potential obstacles to goal satisfaction may also increase the accessibility of goals to which they are related. For example, Fishbach *et al.* (2003) examined the goal of getting good grades in college students. They found that words for temptations that would take people away from studying (e.g. television) increased the accessibility of words relating to studying (e.g. homework), but that the reverse did not occur. Furthermore, this effect was strongest for those who reported themselves as being good at keeping up with their schoolwork. Thus, effective self-regulation is aided when items that might hinder the satisfaction of a goal increase the accessibility of the threatened goal.

In this section, we have given a brief overview of data on the relationship between the activation of goals and the accessibility of concepts related to those goals. This work suggests that goals can activate means for achieving the goal. Means can increase the accessibility of unfulfilled goals. Goal fulfilment can inhibit means. Furthermore, temptations that would block goal satisfaction can increase the accessibility of the goals that are threatened. Furthermore, there is evidence that these changes in accessibility influence cognitive performance. Increasing the accessibility improves performance on goal-related tasks (presumably by allowing people to devote more cognitive resources to the task and by making goal-related concepts easier to use). In addition, changes in accessibility are mirrored by changes in preferences, which can help drive people toward objects that will help them to satisfy unfulfilled goals. In the next section, we turn to influences of motivation on learning.

3.2. *Motivation and learning*

Perhaps it is not surprising that goals influence planning, preferences, and action. After all, for an agent to survive, it must take actions to achieve desirable end states and to avoid undesirable ones. Any movement toward achieving a desirable outcome will be seen as a preference for that outcome over others.

Based only on findings like those described in the previous section, one might conclude that cognitive science need not be that concerned with studies of motivation, except for models of action, decision making, and planning. Models of those processes have typically taken goals into account, and so new research on the Psychology of goals can be expected to be useful.

As we discuss in this section, however, the influences of goals on information processing are pervasive, and influence processes that go well beyond obvious cases of action, and planning. To demonstrate this point, we briefly review two sets of studies that we have conducted exploring the relationship between motivation and a simple perceptual learning task.

3.3. *Regulatory fit and perceptual classification*

Much research in cognitive psychology has explored people's ability to learn simple categories in laboratory settings (e.g. Medin and Schaffer 1978, Nosofsky 1986, Maddox and Ashby 1993, Markman and Ross 2003). This work has provided significant insight into low level learning processes underlying the acquisition of mental representations. In a very simple task, the items might be dots on a computer screen. As shown in figure 1, the categories may overlap so that there are regions of the screen that are ambiguous where the dots might come from either of two categories. On each trial people are shown an item and are asked to classify it into one of a small number of categories. After making a response, they are usually given feedback about whether the response was correct, and then another item is presented. Learning typically continues for a fixed number of trials or until some performance criterion is reached.

Research on classification typically does not take motivation into account. Participants are usually given a simple explicit goal to learn the categories to the best of their ability. The feedback consists of a mild reward (positive verbal feedback or a small monetary reward for exceeding a criterion level of performance). If motivational factors like the ones we have discussed in this paper so far influence

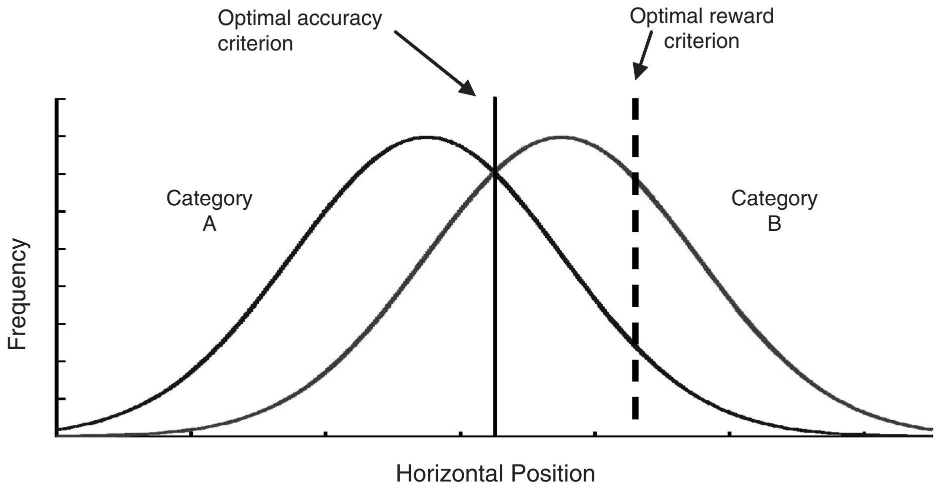

Figure 1. Category distributions with the optimal decision criteria (source: Markman *et al.* in press).

category learning, then it would suggest that the processes that construct mental representations are deeply influenced by an individual's motivational state.

In one study addressing this issue, we had people perform a very simple classification task in which the materials were dots that appeared on a computer screen as shown in figure 1 (Markman *et al.* in press). For each learning task people performed, only one dimension (horizontal or vertical) was relevant to this classification, and the position of the dots that were presented varied only on this dimension. On each trial, people saw a dot and were asked to press a button to indicate which of two categories it belonged to. As in figure 1, the categories were defined by two distributions of dots with means separated in space. The distributions were overlapping, so that some positions on the screen were ambiguous, because dots from either category might appear in those locations. Thus, perfect performance was impossible.

To explore the role of regulatory focus on learning, we manipulated two key aspects of the typical classification task. First, we manipulated regulatory focus. As mentioned above, a situation can give people either a promotion focus (if they are focused on potential gains) or a prevention focus (if they are focused on potential losses). In our study, half of the people were given a promotion focus and half were given a prevention focus. People given a promotion focus were told that they would participate in three classification tasks during the experiment. For each one, if their performance exceeded a criterion, they would receive a raffle ticket for a drawing to win $50. People given a prevention focus were awarded three raffle tickets at the beginning of the experiment. They were told that for each classification task, if their performance exceeded a criterion, they could keep one of the tickets they were given, but if their performance fell below the criterion, they would lose the ticket. Objectively, of course, all participants were in the same situation. They could end the experiment with at most three tickets based on their performance. What differed was only whether participants believed they were achieving tickets with their performance (promotion) or avoiding losing the tickets (prevention).

Table 1. Payoff matrices and performance criteria for the three payoff conditions.

Matrix	Correct High payoff	Incorrect High payoff	Correct Low payoff	Incorrect Low payoff	Initial value	Performance criterion
Mixed	200	−100	0	−100	0	3700
Gain	400	100	200	100	0	33 700
Loss	−111	−411	−311	−411	0	−43 000

In addition to regulatory focus, we manipulated the reward structure of the task. People performed three classification tasks, one using each of the three payoff matrices shown in table 1. Each payoff matrix was asymmetric. That is, one category was associated with a better payoff for correct responses than the other was. With a payoff matrix like this, ideal performance requires biasing responses toward the high-payoff category, because that will maximize payoff in the long run. With this kind of payoff matrix, the response criterion for deciding whether an item is in one category or the other that maximizes long-run reward is different from the one that maximizes long-run accuracy.

As shown in table 1, one payoff matrix had all gains. That is, on every trial, participants received points. Subjects started with zero points and had to achieve the total number of points shown in the table. They received more points for correct responses than for incorrect responses, and more points for correct responses of the high payoff category than of the low payoff category. A second payoff matrix had all losses. In this case, they also started with zero points, but had to keep their total above the negative criterion shown in the figure. They lost some number of points for each incorrect response, but lost fewer points for correct responses than for incorrect responses. Finally, a third matrix had gains for correct responses and losses for incorrect responses. In each task, participants were told how many points they had to have at the end of the block in order to achieve the criterion level of performance.

In this study, we were particularly interested in whether there would be an effect of regulatory fit. As discussed above, regulatory fit effects occur when people's performance is better if their regulatory focus matches the structure of the environment. That is, when people have a promotion focus, they should be most sensitive to gains in the environment. When the payoff structure is all gains, then the environment matches the focus. When people have a prevention focus, they should be most sensitive to losses in the environment. Thus, the payoff structure with all losses should match a prevention focus.

In this task, then, regulatory fit would be observed if people with a promotion focus were better able to maximize long-run reward when they are given a task with a positive payoff matrix than with a negative payoff matrix. People with a prevention focus should be better able to maximize long-run reward when they are given a task with a negative payoff matrix than a task with a positive payoff matrix.

In fact, these results are obtained. An advantage of using these simple classification tasks, is that mathematical models can be used to assess the criterion people are using to respond whether an item is in one category or the other. As regulatory fit would predict, people given a promotion focus had a criterion closer to the one that maximizes long-run reward when they were given the positive matrix than when they saw either the negative or mixed matrix. In contrast, people given a prevention focus had a criterion closer to the one that maximizes long-run reward when they were given the negative matrix than when they were given either the positive or mixed matrix. Thus, people's regulatory focus influenced the nature of the categories people learned.

3.4. *Regulatory focus and flexibility*

The study just described uses a very simple classification task involving categories that are distinguished by only a single relevant dimension. We have begun to explore multidimensional category structures to better understand how regulatory fit influences learning. Our preliminary work suggests that regulatory fit may influence people's willingness to be flexible during category learning, which in turn may affect people's mental representations (Maddox *et al.* in press).

To test this possibility, we created a multidimensional category structure using simple stimuli consisting of lines that could vary their position on a computer screen as well as their length and orientation. Preliminary testing suggested that the position of the line on the screen was most salient for people. We designed the materials so that optimal accuracy with a rule based on any one of the dimensions alone was 83%. However, 100% accuracy could be obtained using a rule that considered both length and orientation. We then gave people a learning criterion of 90%, so that if people used a one-dimensional classification rule, their performance would not exceed the criterion.

We ran two versions of this study. In one, the payoffs were always positive (that is people always gained points on each trial), and in the second, the payoffs were always negative (that is, people lost points on each trial). The payoff matrices were symmetrical, though, so there was no need for people to bias their responses toward one category or another. As before, however, we manipulated regulatory focus. As before, the promotion focus involved telling people they would receive a ticket for an entry into a drawing to win $50 if their performance exceeded the criterion. The prevention focus involved giving people a ticket and telling them that they would lose the ticket unless their performance exceeded the criterion.

Of interest with this design is the speed with which people adopt a categorization rule involving the length and orientation dimensions. We can assess people's rule use by fitting mathematical models to their performance (see Maddox *et al.*, in press, for details). Across the two studies, we found that regulatory fit influences people's flexibility in this task. When the payoff matrix had all gains, people adopted the two-dimensional classification rule involving length and orientation earlier when they had a promotion focus than when they had a prevention focus. In contrast, when the payoff matrix had all losses, people adopted the two-dimensional classification rule earlier when they had a prevention focus than when they had a promotion focus.

These data suggest that motivation can influence mental representations by affecting the strategies that people use during learning. People's cognitive processing is more flexible when their regulatory focus matches the task environment. Thus, we would expect people's mental representations to be richer and perhaps to be a more veridical representation of the outside world when there is a regulatory fit than when there is a regulatory mismatch. This finding is particularly important, because most research on classification has not examined the influence of regulatory focus on learning. Furthermore, most studies use a mild promotion manipulation in which people get positive feedback for correct answers. There is evidence that people differ in their chronic regulatory focus (Higgins 1987), and thus most studies combine people with regulatory fits and regulatory mismatches in the design of the study.

4. Implications for the cognitive architecture

As psychologists, our primary concern is understanding the human cognitive architecture. Thus, the fact that humans are influenced by goal strength and regulatory focus is sufficient for us to continue exploring these factors. In the next section of this paper, we outline some directions for future research in psychology. In this section, however, we discuss the implications of current research on motivation for the design of computational systems.

Two aspects of the motivational system seem like good candidates for inclusion in computational models of mind. First, the data described here suggest that goal activation increases the accessibility of goal-related concepts. For example, means prime goals, and goals prime means. Furthermore, successful self-regulation in humans is facilitated by having potential temptations in the environment activate the goals with which they compete rather than activating the goal for the temptation itself.

These mechanisms are likely to be quite useful for helping computational systems stay focused on information relevant to satisfying an active goal. Philosophical discussions of holism point out that there are significant dangers associated with conceptual systems in which the meaning of one representational element is determined by virtue of its connection to other elements (e.g. Fodor and Lepore 1992). In particular, it can be difficult to determine what information is relevant to a current situation when all mental representations are interconnected, as they are in semantic network models. Indeed, many AI models use representational elements whose meaning is determined in part by their connection to other representational elements.

Semantic network models avoid the problem of holism by having a limited amount of activation flowing through them at any moment, where only those memory elements that have some activation are available to influence cognitive processing (e.g. Anderson *et al*. 1996, Markman 1999). If goals and means are part of that semantic network, then activation of a goals or means will ensure that other information related to that goal will also be accessible. This activation can support opportunistic planning, in which objects in the environment that can be used to satisfy an active goal are perceived and incorporated into action plans (Patalano and Seifert 1997).

A second aspect of the motivational system that has implications for computational models is the distinction between approach and avoidance and the concept of regulatory focus. While some models have incorporated goals into mental representations (e.g. Newell 1990, Anderson 1993), they have not incorporated an explicit distinction between approach and avoidance motivational systems. Because most models are not designed to operate autonomously in real environments and in real time, it has not been necessary to make a strong distinction in computational models between approach goals and avoidance goals. Avoidance goals are particularly important to have when an agent requires mechanisms for avoiding dangerous situations.

The existing data suggest that people may often perform sub-optimally when the environment is not matched to their regulatory focus. In the studies of perceptual classification described above, people given a promotion focus and those given a prevention focus are objectively in the same situation (they will have a ticket for the raffle if their performance exceeds the criterion). Furthermore, the positive, negative and mixed matrices are identical except for whether people are being given rewards are punishments. Thus, any differences in performance caused by regulatory focus are causing suboptimalities in behaviour.

The potential advantage to having distinct regulatory foci is that the cognitive system of an agent can be tuned sensitively to potential gains or losses. Because negative events can happen quickly in potentially dangerous environments, it is useful to be able to set the cognitive system quickly for sensitivity to negative events. Furthermore, promotion focus may give precedence to approach goals, and prevention focus may give precedence to avoidance goals (Markman and Brendl 2005). Thus, regulatory focus influences both the way information is learned as well as the priorities for the types of goals that an agent will pursue. In natural environments, this differential sensitivity is likely to be useful. Future computational research should explore the value of differential regulatory foci for goal pursuit.

5. Future directions for psychological research

We see a number of important avenues for future psychological research to explore the implications of the work described here. First, there are a number of implications of regulatory fit for learning that have not yet been explored. For example, the materials in the studies we have done so far have all been simple perceptual items. However, the content of the categories might also influence regulatory fit. For example, learning about categories of desirable items might be easier with a promotion focus than with a prevention focus, while the opposite might be true for categories of undesirable items.

In addition, categories are used in many different ways, and classification is only one of them. People make predictive inferences with categories, communicate about them, form preferences for category members, and do causal reasoning with categories. Each of these functions may influence what is learned about categories (see Markman and Ross 2003). Regulatory focus may interact with these category uses in interesting ways.

Finally, there has been much recent interest in the relationship between computational models of cognition and the implementation of cognitive systems in the brain. Studies in cognitive neuroscience have used brain-damaged patients and brain imaging techniques in healthy people to study how the brain implements cognitive processes. Research on motivational processes may be useful for extending current work in cognitive neuroscience. Some progress has been made on this front. For example, social and personality psychologists are beginning to team up with computational neuroscientists to develop neurobiologically inspired models of the motivation–cognition interface (Ashby *et al.* 1999, Pickering and Gray 2001).

Similarly, Maddox *et al*. (in press) review an emerging body of literature addressing the neurobiological overlap between motivational systems (such as regulatory focus theory) and classification learning systems whose neurobiology is fairly well understood (Ashby *et al*. 1998, Maddox and Ashby 2004).

6. Conclusions

In this paper, we have summarized recent advances in the interface between motivation and cognition. This work has been driven by experimental techniques that explore the operation of the motivational system indirectly by examining its influence on changes in accessibility of concepts, preferences for items, and speed and quality of category learning. This work has not yet been incorporated into computational models in cognitive science, because these models do not typically focus on systems that act in real environments (Markman and Dietrich 2000b). We view the growing body of work on motivation in psychology as an invitation to cognitive modellers to address the role of motivation in computational models.

References

H. Aarts, P.M. Gollwitzer and R.R. Hassin, ''Goal contagion: perceiving is for pursuing'', *Journal of Personality and Social Psychology*, 87, pp. 23–37, 2004.

J.R. Anderson, *The Architecture of Cognition*, Cambridge, MA: Harvard University Press, 1983.

J.R. Anderson, *Rules of the Mind*, Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1993.

J.R. Anderson, L.M. Reder and C. Lebiere, ''Working memory: activation limitations on retrieval'', *Cognitive Psychology*, 30, pp. 221–256, 1996.

F.G. Ashby, L.A. Alfonso-Reese, A.U. Turken and E.M. Waldron, ''A neuropsychological theory of multiple systems in category learning'', *Psychological Review*, 105, pp. 442–481, 1998.

F.G. Ashby, A.M. Isen and A.U. Turken, ''A neuropsychological theory of positive affect and its influence on cognition'', *Psychological Review*, 106, pp. 529–550, 1999.

J.A. Bargh and T.L. Chartrand, ''The unbearable automaticity of being'', *American Psychologist*, 54, pp. 462–479, 1999.

L.W. Barsalou, ''Perceptual symbol systems'', *Behavioral and Brain Sciences*, 22, pp. 577–660, 1999.

C.S. Carver and M.F. Scheier, *On the self-regulation of behavior*, New York: Cambridge University Press, 1998.

M.S. Fanselow, ''Neural organization of the defensive behavior system responsible for fear'', *Psychonomic Bulletin and Review*, 1, pp. 429–438, 1995.

R.E. Fikes and N.J. Nilsson, ''STRIPS: a new approach to the application of theorem-proving to problem-solving'', *Artificial Intelligence*, 2, pp. 189–208, 1971.

A. Fishbach, R.S. Friedman and A.W. Kruglanski, ''Leading us not into temptation: momentary allurements elicit overriding goal activation'', *Journal of Personality and Social Psychology*, 84, pp. 296–309, 2003.

J. Fodor and E. Lepore, *Holism: A Shoppers Guide*, Cambridge, MA: Basil Blackwell, Inc., 1992.

J. Forster, N. Liberman and E.T. Higgins, ''Accessibility from active and fulfilled goals'', *Journal of Experimental Social Psychology*, 41, pp. 220–239, 2005.

W.S. Geisler and K.L. Chou, ''Separation of low-level and high-level factors in complex tasks: visual search'', *Psychological Review*, 102, pp. 356–378, 1995.

D. Gentner, K.J. Holyoak and B. Kokinov (Eds). *The Analogical Mind*, Cambridge, MA: The MIT Press, 2001.

A.M. Glenberg, ''What memory is for'', *Behavioral and Brain Sciences*, 20, pp. 1–55, 1997.

E.T. Higgins, ''Self-discrepancy: a theory relating self and affect'', *Psychological Review*, 94, pp. 319–340, 1987.

E.T. Higgins, ''Beyond pleasure and pain'', *American Psychologist*, 52, pp. 1280–1300, 1997.

E.T. Higgins, ''Making a good decision: value from fit'', *American Psychologist*, 55, pp. 1217–1230, 2000.

E.T. Higgins, L. Chen Idson, A.L. Freitas, S. Spiegel and D.C. Molden, "Transfer of value from fit", *Journal of Personality and Social Psychology*, 84, pp. 1140–1153, 2003.

L.L. Jacoby, D.S. Lindsay and S. Hessels, "Item-specific control of automatic processes: stroop process dissociations", *Psychonomic Bulletin and Review*, 10, pp. 638–644, 2003.

A.W Kruglanski, J.Y. Shah, A. Fishbach, R. Friedman, W.Y. Chun and D. Sleeth-Keppler, "A theory of goal systems", *Advances in Experimental Social Psychology*, 34, pp. 331–378, 2002.

W.T. Maddox and F.G. Ashby, "Comparing decision bound and exemplar models of categorization", *Perception and Psychophysics*, 53, pp. 49–70, 1993.

W.T. Maddox and F.G. Ashby, "Dissociating explicit and procedure-learning based systems of perceptual category learning", *Behavioral Processes*, 66, pp. 309–332, 2004.

W.T. Maddox, G.C. Baldwin and A.B. Markman, "Regulatory focus effects on cognitive flexibility in rule-based classification learning", *Memory and Cognition*, in press.

W.T. Maddox, A.B. Markman and G.C. Baldwin, "Using classification to understand the motivation–learning interface", *Psychology of Learning and Motivation*, in press.

A.B. Markman, *Knowledge Representation*, Mahwah, NJ: Lawrence Erlbaum Associates, 1999.

A.B. Markman and C.M. Brendl, "The influence of goals on value and choice", in *The Psychology of Learning and Motivation*, D.L. Medin, Ed., Vol. 39, San Diego, CA: Academic Press, 2000, pp. 97–129.

A.B. Markman and C.M. Brendl, "Goals, policies, preferences, and actions", in *Applying Social Cognition to Consumer-focused Strategy*, F.R. Kardes, P.M. Herr and J. Nantel, Eds, Mahwah, NJ: Lawrence Erlbaum Associates, pp. 183–200, 2005.

A.B. Markman and E. Dietrich, "Extending the classical view of representation", *Trends in Cognitive Sciences*, 4, pp. 470–475, 2000a.

A.B. Markman and E. Dietrich, "In defense of representation", *Cognitive Psychology*, 40, pp. 138–171, 2000b.

A.B. Markman and B.H. Ross, "Category use and category learning", *Psychological Bulletin*, 129, pp. 592–613, 2003.

A.B. Markman, G.C. Baldwin and W.T. Maddox, "The interaction of payoff structure and regulatory focus in classification", *Psychological Science*, in press.

A.B. Markman, K. Kim and C.M. Brendl, "The influence of goal activation on preference", in preparation.

D.L. Medin and M.M. Schaffer, "Context theory of classification", *Psychological Review*, 85, pp. 207–238, 1978.

G.A. Miller, E. Galanter and K.H. Pribram, *Plans and the Structure of Behavior*, New York: Holt, Reinhart and Winston, Inc., 1960.

N.E. Miller, "Liberalization of basic S-R concepts: extensions to conflict behavior, motivation, and social learning", in *Psychology: A study of a Science. General and Systematic Formulations, Learning, and Special Processes*, S. Koch, Ed., Vol. 2, New York: McGraw Hill, 1959, pp. 196–292.

A. Newell, *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press, 1990.

R.E. Nisbett and T.D. Wilson, "Telling more than we can know: verbal reports on mental processes", *Psychological Review*, 87, pp. 231–259, 1977.

R.M. Nosofsky, "Attention, similarity and the identification-categorization relationship", *Journal of Experimental Psychology: General*, 115, pp. 39–57, 1986.

A.L. Patalano and C.M. Seifert, "Opportunistic planning: being reminded of pending goals", *Cognitive Psychology*, 34, pp. 1–36, 1997.

A.D. Pickering and J.A. Gray, "Dopamine, appetitive reinforcement and the neuropsychology of human learning: an individual differences approach", in *Advances in Individual Differences Research*, A. Eliasz and A. Angleitner, Eds, Lengerich, Germany: PABST Science Publishers, 2001, pp. 113–149.

R.F. Port and T. Van Gelder (Eds). *Mind as Motion*, Cambridge, MA: The MIT Press, 1995.

J.J. Prinz, *Furnishing the Mind*, Cambridge, MA: The MIT Press, 2002.

J.Y. Shah and A.W. Kruglanski, "When opportunity knocks: bottom-up priming of goals by means and its effect on self-regulation", *Journal of Personality and Social Psychology*, 84, pp. 1109–1122, 2003.

J.R. Stroop, "Studies of interference in serial verbal reactions", *Journal of Experimental Psychology*, 18, pp. 643–662, 1935.

E. Thelen and L.B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge, MA: The MIT Press, 1994.