Prior to publication, this paper should be cited as: (forthcoming) "Why Emotions Do Not Solve the Frame Problem" in V. C. Müller (Ed.), Fundamental Issues of Artificial Intelligence (Synthese Library). Berlin: Springer.
Refer to published paper for the correct pagination, or email me at madeleineransom@gmail.com

Why Emotions Do Not Solve the Frame Problem

By Madeleine Ransom

Abstract: Attempts to engineer a generally intelligent artificial agent have yet to meet with success, largely due to the (intercontext) frame problem. Given that humans are able to solve this problem on a daily basis, one strategy for making progress in AI is to look for disanalogies between humans and computers that might account for the difference. It has become popular to appeal to the emotions as the means by which the frame problem is solved in human agents. The purpose of this paper is to evaluate the tenability of this proposal, with a primary focus on Dylan Evans' search hypothesis and Antonio Damasio's somatic marker hypothesis. I will argue that while the emotions plausibly help solve the intracontext frame problem, they do not function to solve or help solve the intercontext frame problem, as they are themselves subject to contextual variability.

1. What is the frame problem?

The frame problem began as an issue in classical artificial intelligence concerning how to represent in formal logic the effects of actions without having to cumbersomely represent the non-effects of actions (McCarthy & Hayes 1969). When an agent acts, the world changes in some ways, but in many others it stays the same. How can a system update its database (or 'beliefs') to reflect these changes? If one simply excludes the non-effects from the program and only represents the effects of the actions then the problem is that it is not a matter of logic that everything else does in fact stay the same. Though it may be a matter of common sense, the artificial system cannot make this deductive inference from the limited information it

possesses. If one then opts to include the non-effects in the program, the problem is that this quickly becomes computationally intractable, because the number of non-effects one must include is staggeringly large and leads to a combinatorial explosion.

While there are now several adequate methods for addressing this problem in logic-based AI,[1] the frame problem can be expanded into a broader epistemological problem, as Dennett (1978) and Fodor (1983) first noted. Exactly how to characterize this broader problem has remained controversial, however.[2] The issue is that 'the' frame problem is rather a cluster of related problems, though their common core is that each is concerned with how to determine relevance. Wheeler (2008) usefully distinguishes between the *intercontext* and the *intracontext* frame problem. The intracontext frame problem is how to determine, given a context, what information one ought to bring to bear on it – what stored knowledge is relevant in determining what to do? The intracontext frame problem can thus be subdivided into two problems: i) of the many possible actions, which are relevant and thus deserving of consideration? ii) of the many possible consequences of these actions, which are relevant? Computational intractability threatens, given the sheer number of possible actions and consequences available for consideration.

Proponents of simple heuristics hold that the intracontext frame problem is solved by human agents via simple, formalizable rules of thumb that bypass complex search procedures and return an answer that is 'good enough' in a wide enough range of situations. The price they pay is optimality – heuristics don't turn up the best solution in

---

1 See Shanahan (1997). Note, however, that these solutions work only in narrowly fixed domains, and so fail to contribute towards the development of an artificial agent with human scale general intelligence. In the terms I set out below, these proposals can be thought of as solutions to the intracontext, but not the intercontext frame problem.

2 For an overview of the controversy see Ford and Pylyshyn (1996). Gabbay and Woods (2003, 110-1) identify the frame problem with what AI theorists call the relevance problem.

all cases. What they gain is speed and computational tractability, and so the intracontext frame problem is disarmed.[3]

While the exact nature of the solution may be disputed, the general form it will take is agreed upon. All that is needed to solve the intracontext frame problem is positing that we humans possess a handy bag of tricks or shortcuts for navigating our way through the vast stores of information at our disposal. However, even if there is a computationally feasible proposal for how we efficiently solve specific types of problems, a more basic and difficult issue remains – how is it that we are able to know which problem we are facing in the first place? The issue is somewhat analogous to the difficulties facing the student of mathematics who has memorized many formulas, but must then recognize not only when a real life situation calls for a mathematical solution, but also which of these formulas to apply.

The intercontext frame problem is precisely that of how to determine what context one is in. It is the problem of determining what features of the environment one ought to take as relevant, amongst the many possible candidates, in specifying the situation. What makes the problem so difficult is first that such relevance is largely context-sensitive, and second that we are confronted with ever-changing contexts. For example, whether the fact that there is no milk in the fridge is relevant to an agent depends on the context – is she at the supermarket; does she have money; is she about to go on vacation? Even supposing she is at the supermarket, if an earthquake occurs while she's there, or she realizes that she's late meeting a friend for

[3] In another, perhaps compatible approach (see Carruthers 2007), those who hold that the mind has a massively modular architecture may be able to sidestep the intracontext frame problem. These modules are characterized by their dedicated functions – they have more or less narrow purposes and are informationally encapsulated, in that the amount of information the modules are able to draw upon is severely limited. Drastically limiting the amount of information available may render the computational process tractable.

dinner, then the fact that there's no milk in the fridge ceases to be relevant to the situation. The sheer number of potential contexts dashes any hope of specifying a tractable set of rules for determining relevance in all but the narrowest of domains. Dreyfus (1992) characterizes the problem in terms a 'regress of contexts': the present context can only be recognized in terms of features taken to be relevant in a broader context. This broader context can only be recognized in terms of features taken to be relevant in a still broader context, and so on.[4]

A final point on the scope of the frame problem is that it besieges not only practical but also theoretical reason. Not only are we regularly called upon to come up with a swift answer to the question 'what should I do?' but also to that of 'what should I believe?'[5] The intercontext frame problem here is: what evidence is relevant in determining what to believe? The intracontext frame problem is: given the evidence, what information in one's stored database should one draw upon in order to form one's beliefs? That the frame problem extends to theoretical reason is particularly important when evaluating the role of the emotions, given that many discussions have been confined to the domain of practical reasoning.

---

[4] Is the regress infinite? While it may be contexts all the way down, presumably at some point one hits rock bottom. However, while it may not be infinite, the regress is nevertheless still vicious as long as it renders determining context computationally intractable or overly complex and so impractically time-consuming. For example, it would presumably take an extraordinarily long time to determine one's present context if one had to factor in all previous contexts. Stipulating that one only draw upon the relevant factors from previous contexts simply causes the intercontext frame problem to arise anew, as now one must explain how such relevance is determined. Thanks to Chris Mole for discussion on this point.

[5] It is consideration of the second question that has led some philosophers to draw parallels between the problem of induction and the frame problem, though the reduction of the latter to the former is a controversial (and in my view misguided) move. See especially the exchange between Fetzer (1991) and Hayes (1991); and Dennett (1978, 1998).

2. Emotions and the Frame Problem

Attempts to engineer a generally intelligent artificial agent have yet to meet with success, largely due to the (intercontext) frame problem. Given that humans are able to solve this problem on a daily basis, one strategy for making progress in AI is to look for disanalogies between humans and computers that might account for the difference. It has become popular to appeal to the emotions as the means by which the frame problem is solved in human agents. While Herbert Simon (1967) pioneered the integration of affect with human cognition, his suggestion that the emotions serve as 'interrupt systems' fell short of addressing the frame problem head on. Ronald de Sousa first picks up the thread (1979) and then makes the explicit connection (1987), arguing that the "Emotions spare us the paralysis potentially introduced by this predicament [of having to first retrieve information in order to determine whether it is relevant] by controlling the salience of features of perception and reasoning" (172).

In what might be construed as a buildup of momentum, several authors have made claims in recent years that emotions (help) solve the frame problem in human, and perhaps artificial, agents. In her landmark book *Affective Computing* (1997), Rosalind Picard argues – while discussing the problem of combinatorial explosion – that "AI has ignored a crucial component [of human intelligence] that is even more basic to human problem solving abilities: the use of feelings and intuition to guide reasoning and decision making […] An integral component of human decision making is emotion, and this component could potentially be given to computers" (221-2). Dennett (1998) endorses de Sousa's (1979) proposal as a promising avenue for addressing the frame problem, though he laments the lack of a sufficiently concrete scheme for its implementation. Megill & Cogburn (2005) are committed to the hypothesis that "the emotions play a prominent role in preventing humans from suffering from the

frame problem" (311). Ketelaar & Todd (2001) argue that "emotions can help the computationally limited human mind to circumvent the pitfalls of the frame problem by determining *which* information to attend to in the first place" (204, emphasis original). Evans (2004) eschews talk of the frame problem altogether in order to avoid controversy over what the problem really is. However, what he calls the 'search problem' – how to cap the number of consequences of an action under consideration – is equivalent to the second part of the intracontext frame problem. Evans takes himself to be elaborating on de Sousa's view, which he characterizes as that the "non-rational procedure for delimiting the range of consequences to be considered in a rational decision process is governed by the emotions" (181).

Damasio's (1994) somatic marker hypothesis (SMH) is often invoked as the means by which the emotions solve or help solve the frame problem. While Damasio himself does not make explicit reference to the frame problem, he holds that the emotions, in the form of somatic markers, "assist the deliberation by highlighting some options (either dangerous or favorable) and eliminating them rapidly from subsequent consideration" (174).

Just how strong the role proposed for the emotions varies, though this is not always made explicit in the literature. This is in part because it is not always clear what the frame problem is taken to be, and thus which problem the emotions are supposed to be solving. Leaving aside the specifics of just how the emotions are taken to (help) solve the problem in each case, claims can be mapped onto the following taxonomy:

H1: the emotions help solve the intracontext frame problem
H2: the emotions help solve the intercontext frame problem
H1*: the emotions solve the intracontext frame problem
H2*: the emotions solve the intercontext frame problem

One can also be committed to weaker or stronger versions of H1 or H2. Recall that the intracontext frame problem consists of two subdivisions: i) of the many possible actions, which are relevant and thus deserving of consideration? ii) of the many possible consequences of these actions, which are relevant? So a weaker version of H1 might hold that the emotions help only with (i) but not (ii), a stronger version that the emotions help with both (i) and (ii). The specific nature and extent of the help the emotions offer will also strengthen or weaken H1 and H2.[6] A further distinction amongst hypotheses that can be made is whether the emotions are taken to (help) solve the (intra/inter-context) frame problem in the domain of practical reasoning only or also in that of theoretical reason. Even if one holds that the emotions alone solve the intracontext frame problem in the domain of practical reason, if this solution does not extend to theoretical reason then the strongest claim one is committed to is H1.

3. The somatic marker hypothesis

Damasio (1994) first proposed the SMH as a way of accounting for the behavioral anomalies of numerous patients with damage to the ventromedial prefrontal cortex (VMPFC). While the means-ends reasoning skills of these patients remained unaffected, along with a host of other mental capabilities, they exhibited a striking inability to

6 Does the strength of the hypotheses also depend on what one counts as an emotion? Perhaps the broader and more inclusive the class or natural kind, the more resources one has at one's disposal for solving the frame problem. On the other hand, increasingly complex and cognitive emotions seem especially subject to the criticisms I make in this paper. However, even if only some subset of the emotions are implicated in solving the frame problem, one could still adhere to H1* or H2*, if these select emotions solve the frame problem on their own. Thanks to Adam Morton for the question, and see Morton (2013) for an example of an inclusive view of the emotions.

make wise – or sometimes any – practical decisions regarding their own lives.

The SMH holds that the emotions are involved in practical reasoning – they help us to make decisions concerning what to do.[7] Emotional feelings serve as somatic markers. They 'mark' the content of mental states with either a positive or negative emotional valence.[8] Such markers may then be reconstituted by the agent during deliberation, and thus help guide behavior.

However, as Linquist and Bartol (2012) have pointed out, the SMH isn't just one, but actually a series of separable hypotheses concerning how the emotions are involved in practical reasoning. There are at least five conceptually separable stages in the decision making process where somatic markers may play a role.[9]

In the first stage, decision point recognition, a reconstituted somatic marker functions to alert the agent that a decision should be made. The second stage is that of generating candidate options – somatic markers work to heighten, or 'energize' working memory and attention, allowing the options to be generated. In the third stage, deliberation, the agent identifies the implications or additional properties of various options. The way somatic markers are implicated in the pro-

---

[7] Linquist & Bartol (2013) make a distinction between the Somatic Marker Hypothesis and the Somatic Marker Model: "The somatic marker model…describes a putative neurocognitive mechanism for associating autonomic tags with mental representations. Somatic marker hypotheses, in contrast, invoke this model to explain some aspect of cognition, such as practical decision making" (458). So, strictly speaking, the SMH may hold that somatic markers are employed in processes other than practical reasoning, such as theoretical reasoning.

[8] Damasio here appears to be using valence to mean what Colombetti (2005) terms 'affect' valence: how good or bad an emotion feels. As she points out, however, the term is used in multiple and sometimes conflicting ways in the literature on the emotions.

[9] To be clear, these are not necessarily sequential stages – many of these may occur in parallel and feed into each other.

cess can be sub-divided into two categories. The *relevance hypothesis* is that somatic markers are involved in helping to identity factors relevant to the decision at hand. The *search hypothesis* is that somatic markers are what put a cap on the time and energy we spend deliberating. The fourth stage is that of value assignment and ranking. Value assignment occurs when an option is considered. The somatic marker triggered by the option serves as a factor that weighs for or against it – somatic markers lend valence to the various options. Value ranking is then a means of ordering the various options, thus allowing the option at the top of the hierarchy to emerge as the chosen plan of action. Somatic markers accomplish this by tabulating the valences associated with each option, with the option with the highest overall positive valence winning out. In the fifth stage, somatic markers serve as the motivators for action – once the agent has arrived at a given course of action, somatic markers provide the drive to execute.

4. Why somatic markers don't solve the frame problem

4.1 The SMH and the intercontext frame problem

The intercontext frame problem arises because there are lots of things we might potentially pay attention to in order to determine context. Decision point recognition thus belongs to the intercontext frame problem – recognizing that a decision is called for depends on being able to pick out the features relevant to identifying a (change of) context. The SMH holds that people solve this problem by paying attention to the features of the environment that come with a somatic marker attached. On the stronger reading of this suggestion, corresponding with H2*, people pay attention *only* to those features of the environment that are somatically marked – emotions solve the intercontext frame problem. On the weaker reading, H2, people pay attention *primarily* or *in part* to the valenced features of their sur-

roundings, but this must be supplemented by other strategies or information that guide attention.

The problem with the strong version of the claim – that emotions solve the intercontext frame problem – is that the valence of objects is itself often context dependent. Most objects, people, places, and states of affairs possess both positive and negative aspects, which are often highly variable depending on context. A knife in one context is a helpful tool for making supper and in another it's a threat to one's wellbeing. The prospect of taking a test produces a very different feeling depending on whether one has studied or not. Moreover, there are many concepts which, when combined, elicit an emotional reaction that neither elicits in isolation. As Darwin (1872) observed, a man's beard with some soup caught in it is disgusting even though one typically considers neither soup nor beards to be disgusting. Such contextually elicited emotion is ubiquitous. Which of the multiple somatic markers associated with a given object or situation should be reconstituted in a given situation? Well, the answer goes, it depends on the context. Recognition of which context one is in must therefore come *before* one can employ the appropriate somatic markers, and so they cannot be invoked to solve the intercontext frame problem.

Might the emotions nevertheless *help* solve the intercontext frame problem? Absent a proposal as to what other elements are involved in solving the problem, this claim is difficult to evaluate. However, given that somatic markers themselves depend on context for deployment, the extra component needed to supplement the account appears to be the kind of thing that would itself determine the context, so this other mechanism would be doing all the work. Therefore, the emotions do not solve, nor do they help solve, the intercontext frame problem.

At this point, one might run an objection as follows. Note that H2 itself is ambiguous: it could be either that the emotions solve the intercontext frame problem on their own in a narrow range of cases, or that they are always mere helpers in a wide range of cases; and the first interpretation is not subject to my criticism of H2 here, so it is still a live option.[10] However, then the challenge is how to make the first interpretation tenable – one must explain how it is that the emotions solve the intercontext frame problem on their own in some limited domain. One might proceed by pointing out that while many objects and states of affairs are multi-faceted in their valences, some may be consistently positively or negatively valenced, independent of context. Take the case of fear, for example. When a rabbit spots a hawk swooping towards it, it just runs away, no matter what activity it was engaged in beforehand. The fear the rabbit feels appears to be what initiates the action, with no need to determine context beforehand. So at least some emotions are capable of guiding action while cutting out the middleman of context.

The first thing to notice here is that while hawks are consistently negatively valenced for rabbits, there are presumably few objects like this for humans. Our world is more nuanced and complex than that of a rabbit, and the vast majority of the objects and situations we encounter will be multi-faceted in their valences. So even if there are a few consistently negatively valenced objects for us, it's hard to see how this will take us any significant distance towards solving the intercontext frame problem. If this is the extent of the help the emotions offer us, then it is exceedingly weak indeed.

Secondly, the action that fear initiates (or perhaps only motivates) in the rabbit does not engage practical reason at all; the action is swift and reflex-like. It's hard to see, therefore, how this sort of action program could be a solution to the intercontext frame problem at all.

---

[10] Thanks to Dominic McIver Lopes for this point.

The frame problem arises on the assumption that we are in fact capable of acting intelligently, where intelligence is defined along the lines of the ability to respond in an adaptive manner to ever-changing contexts (Wheeler 2008), and this is presumed to take place through a rational process of sorts, or at least a flexible cognitive process. While rigid, reflex-like actions may account for our ability to leap out of the way of oncoming threats, for example, they are ill-suited to capture the sort of intelligence we take human beings to possess – the capability of responding to new circumstances with the flexibility needed to navigate them in an advantageous manner. And to respond in such a way, the identification of contextual factors is paramount.

4.2 The SMH and the intracontext frame problem

The third and fourth stages of decision-making are where the emotions do the work of addressing the intracontext frame problem. In deliberation, somatic markers flag certain options as relevant, and determine the time to be spent deliberating. In value assignment and ranking, somatic markers serve to further 'prune down' the number of options.

However, for Damasio's proposal to be tenable, he must endorse something along the lines of Newell & Simon's (1976) conception of problem solving as a search through a state space. On this model, rather than generating all the options then pruning them down, individual options are generated and tested step by step.[11] This is precisely the tactic Evans (2004) opts for, and so Damasio's account can be supplemented at this point. Evans takes the emotions to solve

11 The reason Damasio must endorse this model is because first generating all possible options and then sifting through them simply will not help solve the intracontext frame problem – the sheer number of possible options for any given problem will be enormous, and in some cases infinite.

what he calls the search problem, or the problem of when to stop listing the possible consequences of actions.[12] Invoking Newell and Simon's method, the process of searching for a solution to a given problem can be likened to building a search tree. Potential actions represent the first level of nodes on the tree. Their potential consequences represent the second level of nodes, the consequences of those consequences represent a third level, and so on, with the branches becoming ever denser as one expands the tree. While the tree one can build up is in principle infinite, in practice a good search strategy and test delimit the number of branches to be developed. A search strategy determines which node of the tree ought to be expanded first.[13] Whenever a node is expanded a test is then applied to the result to determine whether it constitutes an acceptable solution. On Evans's view, "emotions prevent us from getting lost in endless explorations of potentially infinite search spaces by providing us both with the right kind of test and the right kind of search strategy for each kind of problem we must solve" (185).

Unfortunately Evans provides no account of how the emotions could be employed to determine which node of the tree to expand first. Perhaps they might work in the form of a rule such as 'expand the most emotionally salient option first.'[14] For such a rule to function, we would need to be able to emotionally appraise the options without considering their consequences, for if we did so we would thereby be expanding the nodes on the tree. This rule would be adequate for cases where the somatic markers are tied directly to the options

---

12 The search problem can thus be identified with part (ii) of the intracontext frame problem.

13 One may, for instance, expand all nodes at the first level before moving on to the next level, or one may choose to expand one particular first-level node on the tree to a fixed depth before moving onto other nodes if the first node provides an unsatisfactory solution.

14 What if one is indifferent between the options? Then a supplementary rule might added – indifference is a signal to expand all of the first level nodes to the next level, thus taking their consequences into consideration.

themselves – perhaps we have already contemplated or experienced their consequences and so have come to mark the option directly.

What test do the emotions provide? It appears that on Evans's view, the emotions function as a test for the viability of a given option by providing valence information. Evans defers back to Damasio at this point – the somatic markers associated with certain consequences count for or against this course of action. If their valence is sufficiently negative, then the option is eliminated. If it's sufficiently positive, then the option is chosen.[15]

However, while the emotions may help to solve the intracontext frame problem, they cannot solve it on their own. The first issue is that in many cases the constraints on what options are generated for consideration are not wholly determined by the emotions. Recall that on Damasio's proposal, somatic markers serve only to sustain the option generation process – they are not implicated at all in the formation of the options themselves. Even supposing they are involved, environmental constraints such as the resources one has at one's disposal, along with background knowledge, habit, and other factors will also go into constraining the types of options that are generated in the first place. For example, taking a private helicopter to work is not an option for most of us – not for lack of enjoyment, but rather due to lack of resources. It is for this reason that the strong claim doesn't go through. The emotions cannot solve the intracontext frame problem on their own because they cannot be wholly responsible for constraining what *sorts* of options are generated in the first place on their own. So the first component of the intracontext frame

15 What counts as sufficient? While Evans doesn't provide an account, one may suppose there is some threshold, contextually determined by the importance of the situation. Problems deemed extremely important will perhaps require a higher positive score for a given option to be chosen, or a lower negative score for an option to be eliminated. The emotions may be further implicated here, as the strength of one's feelings may serve as a proxy for the subjective importance of the problem, and so serve to set the threshold itself.

problem – how to select only the relevant possible actions for consideration – is not wholly resolved by the emotions.[16]

The second issue is that not all cases of practical reasoning are likely to involve somatic markers in any significant way. While there are many emotionally charged decisions we must make in life, so too are there many that will leave us cold. These are not cases where we are indifferent as to the outcome, but rather cases where there aren't enough somatic markers associated with the consequences in order to make this method of option elimination or selection useful. Perhaps there is as of yet no somatic marker associated with the relevant consequences. Perhaps the decision simply calls upon background knowledge rather than somatic markers. This is the second reason why the strong claim is untenable – the emotions can at most help solve the intracontext frame problem because they can't do all the work of eliminating or selecting options in all cases.

Further support for this second point comes from the research of Damasio and colleagues on patients with damage to the ventromedial prefrontal cortex. The evidence suggests that the emotions only play a role in reasoning in situations that directly involve the agent.[17] In the lab, such patients successfully navigate the many possible options to generate reasonable solutions to moral, social or instrumental problems. Given that such patients are hypothesized not to be able to reconstitute their somatic markers when reasoning, such markers cannot be necessary for successful performance on these types of tasks. Therefore, the strongest tenable hypothesis is that somatic markers serve to *help* solve the intracontext frame problem.

---

16 Given that Evans's search hypothesis addresses only the second component of the intracontext frame problem, it's unclear that he means to propose a solution to this first issue.

17 See Damasio (1994), part I.

A third issue is that the emotions are unlikely to be of much use in solving problems of theoretical reasoning. If one accepts that theoretical reasoning also runs up against the intracontext frame problem, then a proposal about how somatic markers might be of use here is in order. It's hard to see how one might deliver this, though, especially given that the VMPFC-damaged patients show no theoretical reasoning deficits.[18]

A final and more general problem with invoking the emotions is that there is not always a straightforward connection between avoidance and an object or state of affairs that is marked as negative. We have all carried out actions we know we ought to, in spite of not 'feeling' like it. In addition, we actively seek out some fear-producing items, such as horror movies and public speaking. We also seek to avoid many items that we consider pleasurable, such as cigarettes and junk food. The way we assign value to states of affairs thus appears more complicated than mere emotion, suggesting perhaps that value assignment and ranking is not accomplished solely by somatic markers. This in turn speaks to the fact that the emotions can only help solve the frame problem.

In its most successful incarnation then, Damasio and Evans's proposal will be quite weak: it will take emotions to be kinds of heuristics that work along with other heuristics to shrink the space of possible options. The emotions, on this account, are just one method among many to cut down on the number of actions and consequenc-

18 There has been recent talk of the epistemic emotions serving as heuristic devices, via somatic markers (Hurley et al. 2011). While the workings of the epistemic emotions remains underexplored, I suspect they will be subject to many of the same objections I raise here. An important challenge the proposal faces is to explain how VMPFC-damaged patients manage to perform well on theoretical reasoning tasks. By hypothesis, these patients cannot reconstitute their somatic markers, so it appears the emotions are not necessary to theoretical reasoning. A possibility (thanks to Samantha Matherne) is that the brain-damaged patients do in fact exhibit some sort of limited theoretical reasoning deficit, perhaps in analogical reasoning.

es that need to be considered. So while the proposal can be seen as a serious contender for helping to solve the intracontext frame problem, it falls short of solving the problem on its own. However, one positive element that emerges from this account is that it suggests new directions for empirical research, aimed at exploring the emotions as heuristics model.[19]

5. Conclusion

While the emotions may initially appear to offer a promising solution to the frame problem, their helpfulness is severely limited. The emotions don't solve or help to solve the intercontext frame problem because the valence associated with many objects and states of affairs is itself context dependent. The emotions cannot usefully direct us towards relevant features of our environment, because they in turn rely on those same relevant features for their deployment.

The emotions don't solve the intracontext frame problem on their own for at least four reasons. First, the emotions are not uniquely responsible for selecting only the relevant possible actions for consideration. Second, not all cases of practical reasoning are likely to involve somatic markers in any significant way. Third, they are unlikely to be of much use in solving problems of theoretical reasoning. Fourth, there is not always a straightforward connection between avoidance and an object or state of affairs that is marked as negative. Therefore, the strongest viable claim is the weak hypothesis, H1, that emotions help solve the intracontext frame problem. Given the diminished prospects for resolving the frame problem in human agents via the emotions, it is unlikely that they will be of

---

19 While Ketelaar & Todd are advocates of the heuristics research program, their claim that the emotions solve the frame problem is too ambitious. The proposal for emotions as heuristics here is more modest.

much use in addressing the problem with respect to generally intelligent artificial agents.[20]

20 For helpful discussion and comments on earlier drafts of this paper thanks to Chris Mole, John Woods, the faculty and graduate students of the University of British Columbia, and the audience of the 2013 Philosophy and Theory of AI conference at Oxford – especially Michael Wheeler and Murray Shanahan.

References

Barrett, L.F. (2012). Emotions are Real. *American Psychological Association,* 12(3), 413–429.

Carruthers, P. (2007). Simple Heuristics Meet Massive Modularity. In P. Carruthers, S. Laurence, & S. Stich (Eds), *The Innate Mind: Volume 2: culture and cognition* (pp. 181-198). Oxford: Oxford University Press.

Colombetti, G. (2005). Appraising Valence. *Journal Of Consciousness Studies* 12(8-10), 103-126.

Damasio, A. (1994). *Decartes' Error*. New York: Grosset/Putnam.

Darwin, C. (1872/2009). *The Expression of the Emotions in Man and Animals*. New York: Penguin Classics.

Dennett, D. (1978). *Brainstorms*. Cambridge, Mass.: MIT Press.

Dennett, D. (1998). Cognitive Wheels: The Frame Problem of AI. In *Brainchildren* (pp. 181-205). New York: Penguin Books.

De Sousa, R. (1979). The Rationality of Emotions. *Dialogue,* 18(1), 41-63.

De Sousa, R. (1987). *The Rationality of Emotion*. Cambridge, Mass.: MIT Press.

Dreyfus, H.L. (1992). *What Computers Still Can't Do*. Cambridge, Mass.: MIT Press.

Evans, D. (2004). The Search Hypothesis of Emotion. In D. Evans, & P. Cruse (Eds), *Emotion, Evolution, and Rationality* (pp. 179-192). Oxford: Oxford University Press.

Fetzer, J.H. (1991). The Frame Problem: Artificial Intelligence Meets David Hume. In K.M. Ford, & P.J. Hayes (Eds), *Reasoning Agents in a Dynamic World: The Frame Problem*. Oxford: JAI Press.

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

Ford, K.M. & Pylyshyn, Z.W. (Eds) (1996). *The Robot's Dilemma Revisited: the frame problem in artificial intelligence*. Norwood, NJ: Ablex.

Gabbay, D. & Woods, J. (2003). *Agenda Relevance: a study in formal pragmatics.* New York: North-Holland.

Hayes, P.J. (1991). Artificial Intelligence Meets David Hume: A Reply to Fetzer. In K.M. Ford, & P.J. Hayes (Eds), *Reasoning Agents in a Dynamic World: The Frame Problem*. Oxford: JAI Press.

Hurley, M., Dennett, D. & Adams, R. (2011). *Inside Jokes: Using Humor to Reverse Engineer the Mind*. Cambridge, Mass.: MIT Press.

Ketelaar, T. & Todd, P.M. (2001). Framing our Thoughts: ecological rationality as evolutionary psychology's answer to the frame problem. In P. Davies, & H.R. Holcomb (Eds), *Conceptual Challenges in Evolutionary Psychology: Innovative Research Strategies* (pp. 179-211). Dordrecht: Kluwer Publishers.

Linquist,S. & Bartol, J. (2013). Two Myths about Somatic Markers. *The British Journal for the Philosophy of Science*, 64(3), 455-484.

Megill, J.L. & Cogburn, J. (2005). Easy's Gettin' Harder all the Time: the computational theory and affective states. *Ratio,* 18(3), 306-316.

McCarthy, J. & Hayes, P.J. (1969). Some Philosophical Problems from the Standpoint of Artificial Intelligence. In D.Michie and B.Meltzer (Eds), *Machine Intelligence 4* (463-504). Edinburgh: Edinburgh University Press.

Morton, A. (2013). Emotion and Imagination. Cambridge, Mass.: Polity Press.

Newell, A. & Simon, H.A. (1976). Computer Science as Empirical Inquiry: Symbols and Search. *Communications of the Association for Computing Machinery,* (19), 113-126.

Shanahan, M.P. (1997). *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. Cambridge, Mass.: MIT Press.

Simon, H.A. (1967). Motivational and Emotional Controls of Cognition. *Psychological Review,* 74(1), 29-39.

Wheeler, M. (2008). Cognition in Context: phenomenology, situated robotics and the frame problem. *International Journal of Philosophical Studies*, 16(3), 323-349.