

THE ROLE OF EMOTION IN PROBLEM SOLVING

Roman V Belavkin

School of Computer Science and Information Technology
University of Nottingham, Nottingham, NG8 1BB, UK
rvb@cs.nott.ac.uk

Abstract

Performance and data from some cognitive models suggested that emotions, experienced during problem solving, should be taken into account. Moreover, it is proposed that the cognitive science approach using both theoretical and experimental data may lead to a better understanding of the phenomena. A closer investigation of ACT-R cognitive architecture (Anderson 1993) revealed some properties analogous to phenomena known from the activation theory of emotion. A model of the classical Yerkes-Dodson experiment was built to test the predictions. The study explained such psychological phenomena as arousal, motivation and confidence within the mathematical notation. The influence of changes in these motivational states, controlled by emotion, on information processing has been investigated and it is shown that the dynamics corresponds to the well-known optimisation methods, such as best-first search and simulated annealing.

1 Introduction

Recent progress in cognitive modelling has allowed the testing of quite a broad range of human psychology and cognition theories. A lot of work in experimental psychology has been reproduced and reconsidered by cognitive scientists within theories such as SOAR (Newell 1990) or ACT (Anderson 1993). This work has produced new insights into our understanding of some phenomena of human memory, learning, perception and reasoning. Yet there have been few (if any) attempts to understand emotion and affect within this framework.

The subject of emotion has bothered many psychologists, philosophers and neurobiologists since William James first attempted to define emotion. Many theories of emotion, sometimes quite contradictory, have appeared since then (see (Plutchik 1994, LeDoux 1996) for reviews). Recently the subject of emotion has attracted the attention of the computer science and artificial intelligence communities, and has emerged into a new area of research, sometimes referred as *affective computing* (see (Picard 1997) for review).

Although there is no doubt emotion is a very important component of human and animal psychology, one of the most intriguing and interesting question remains unanswered: Is affect and emotion the necessary component of intelligence? And if it is, how should it be included into AI theory?

It is known from the whole history of experimental psychology that emotion is closely related to cognitive processes such as learning, decision making and memory. Recently there have been claims, based on some experimental evidence, that damage to emotion responsible areas of the brain impairs these cognitive processes (Damasio 1994), and therefore emotion is the necessary

component of intelligence, as also suggested by Goleman (1995). But, as noted in (Sloman 1999), this conclusion is a little premature, as the role of the damaged brain areas is still not sufficiently well understood due to the complexity of the human brain.

Perhaps, in order to support these experimental observations, one needs some theory explaining these phenomena as a mathematical model. We believe that such a theory can be created within the frame of cognitive science. If the role of emotion in cognition can be understood within this theory then, perhaps, we shall also be able to decide what role it plays in intelligence.

It is becoming evident that cognitive models used to test different psychological theories should take emotion into account (Belavkin, Ritter & Elliman 1999). Many of these models simulate subjects solving various puzzles and problems, some models consider children, as in (Jones, Ritter & Wood 2000), whose emotions are easily observable. We know from our own experience and observation that emotion accompanies any problem solving process, but these computer simulations say nothing about it. It seems that in cognitive science this subject has not yet been studied deeply enough.

This is not to claim that cognition and emotion were not considered together at all, but there was not many attempts to study this subject with cognitive architectures such as SOAR or ACT-R. Although there is already a number of cognitive appraisal models, such as (Ortony, Clore & Collins 1988) or (Roseman, Antoniou & Jose 1996), that allow us to conclude which emotion a subject should feel under given circumstances, these symbolic representations do not explain what happens to the thinking process itself as a result of these emotions. What is the difference between assembling the tower of Hanoi in an angry or a happy mood?

2 Modelling the cognition

Although cognitive science is moving towards a unified theory of cognition, there are several schools and architectures implementing different interpretations. Many ideas in this work were developed after observation of some cognitive models, implemented in ACT-R (Anderson 1993), and after closer investigation of the ACT-R conflict resolution mechanism.

2.1 ACT-R architecture and models

The main distinguishing feature of the ACT-R architecture is its sub-symbolic processing capabilities. The mechanism, called *rational analysis*, acts underneath the production system and it can be seen as a stochastic optimisation mechanism. In brief, all the symbols in ACT-R (declarative knowledge units and production rules) have some activation values and associations between them with different strengths. Values of these activations and associations affect many processes in the production system, such as retrieval of knowledge facts, conflict resolution (the choice of one production rule from several satisfying the current condition) and these values are statistically learned, may decay in time and be affected by global parameters.

Such a mechanism can be controlled through many parameters, such as noise variance, goal value, retrieval and utility thresholds, etc. These parameters may dramatically alter the behaviour of the production system and are used by cognitive modellers to adjust the performance of their models and test some theories.

Although ACT-R was sometimes criticised for having too many parameters (and hence a good correlation with the data), there is a great deal of experimental evidence for including them in the theory. There is even a neural implementation ACT-RN (Lebiere & Anderson 1993) (“N” for neural), and the latest version of ACT-R possesses most of its connectionist properties while still retaining the high level of abstraction allowing for encoding and solving complex cognitive problems.

2.2 The Tower of Nottingham

As was mentioned earlier, many ideas for this work were inspired by the results of the Tower of Nottingham model (Jones et al. 2000), which was used to study cognitive development. The idea of this work was to create a model resembling the behaviour of adult subjects assembling the Tower of Nottingham puzzle (Figure 1) and then to achieve a match with the data from seven years old children by modifying the original adults model.

The model was implemented in ACT-R and used the Nottingham “Eye and Hand” perception-action module (Baxter & Ritter 1996) to interact with the task simulation. It achieved a fair match with the data from adult subjects at default parameters settings. Then, to match




Figure 1: The Tower of Nottingham and two of 21 wooden blocks used to build the Tower of Nottingham

young problem solvers, Jones tried several architectural changes using parameters, such as number of chunks in working memory, retrieval threshold, fovea and parafovea sizes (perception) and others.

It is not necessary to describe all the results of that work here, but there was one particular adjustment to the model, that alone produced excellent results (other single parameters could not produce such a good correlation). It was the model with increased noise in conflict resolution. The corresponding parameter :egn (expected gain noise), when set to a value of 6.0, produced a particularly good match for the data such as time needed to complete each layer or the number of constructions assembled on each layer (Figure 2. From (Jones et al. 2000)).




Figure 2: Time taken to assemble each layer and number of constructions assembled on each layer of the Tower by children and EGN6 model.

The fact that children seem to be more “noisy” or, speaking in ACT terms, less rational than adults lead to some interesting speculations and questions. For example, it may indicate that emotions play a greater role in children’s learning and problem solving. Indeed, joy and frustration are more observable in children. Could we possibly model younger children by increasing the noise even further? Or could these or better results be produced by some other parameters in conflict resolution?

3 Decision making in ACT-R

The questions mentioned in previous section were studied by the author during the experiments with the Tower of Nottingham model and a closer investigation of ACT-R conflict resolution mechanism formed the basis for this study. We do not need to explain here the whole spectra of mechanisms and parameters in ACT-R and we shall refer to (Anderson & Lebiere 1998) for further information, but

it is important to give a brief introduction to the ACT-R conflict resolution mechanism and its notation, since it will be used in the rest of the paper. Those familiar with ACT-R may skip the next section.

3.1 Conflict resolution in brief

In ACT conflict resolution (a process of selecting one rule out of several matching the condition) is realised through its *rational analysis* model, which uses the subsymbolic information. Every rule, in addition to its symbolic representation, has also so called *expected gain* E and a rule with the highest gain wins the competition in a conflict set.

Expected gain E is calculated by the following equation:

$$E = PG - C + \xi(\tau), \quad (1)$$

where P is expected probability of achieving the goal if the rule fires, G is the value of the current goal in time units, C is expected cost of that rule in time units (it represents how long will it take to achieve the goal if that rule fires), and $\xi(\tau)$ is a random variable representing noisy or non-deterministic part of ACT-R conflict resolution mechanism. The level of this noise can be controlled through a global variable called *expected gain noise variance* σ and we shall refer to it as the noise temperature $\tau = \sqrt{\sigma}/\pi$.

Although there is a lot of experimental evidence in favour of ACT rational analysis model, unfortunately, the ACT-R book (Anderson & Lebiere 1998) does not give any justification or reference to a source of above formula (perhaps, assuming that every reader is familiar with Bellman's dynamic programming theory). In this paper the probabilistic interpretation of expected gain will be introduced and we shall see how equation (1) can be derived from it. First let us take a closer look at this formula.

As we can see the expected probability P and expected cost C are properties of the production rule. These properties can be learned statistically and there is also a mechanism to "forget" some of this information with time, called probability decay. So, if there are two rules matching the current goal and ACT-R model learned from previous experience that applying the first rule will lead to the goal in 5 minutes ($C = 5$) with expected probability $P = .5$, while for the second rule $C = 10$, $P = .8$ respectively, then for $G = 20$ (default value) the expected gains of these rules will be $E_1 = .5 \cdot 20 - 5 = 5$ and $E_2 = .8 \cdot 20 - 10 = 6$. In this example the second rule will be selected for $E_2 > E_1$.

But not only statistical information about expected probabilities and costs of a rule affect the conflict resolution. The goal value G is a property of the current goal and it is global parameter as well as the noise temperature τ . In the above example one can easily check that for a lower goal values, such as $G = 14$, the first rule will have higher expected gain and it will be selected, despite its lower expected probability. Also, if the noise variance is too high

the choice will become more random and less dependent on P s and C s of rules in conflict set.

As mentioned above, there is a great deal of experimental evidence confirming the plausibility of ACT-R's choice mechanism. For example, the data from (Friedman, Burke, Cole, Keller, Millward & Estes 1964) shows that although subjects choose according to the probability of success (or reinforcement), the proportion of choices of an alternative with maximum success probability ($P = 1.0$) never reaches 1.0. Similarly, subjects still sometimes choose even the most unfortunate alternative (with 0 success probability). So, there is always some degree of randomness (noise) in their choice. Other works on choice probability, such as (Myers, Fort, Katz & Suydam 1963), showed that the choice depends more on the probability of success for higher rewards (goal value).

3.2 Asymptotic properties of rationality

Only one rule can be selected to fire on each cycle and with n production rules in the conflict set, the probability of selecting a particular i -th one is given by Boltzmann equation:

$$p_i = \frac{e^{E_i/\tau}}{\sum_{j=1}^n e^{E_j/\tau}}, \quad (2)$$

where E_i is the evaluation of i -th rule, which is $P_iG - C_i$. We are interested in how the choice probability p_i depends on P_i and C_i for extremely high or low values of goal value G and noise temperature τ .

For simplicity, let us consider the case of two production rules and let P_1 and P_2 be their expected probabilities and C_1 and C_2 their costs. We may also take $P_2 = 1 - P_1$. Then the probability p_1 of choosing the first of two rules will be calculated as

$$p_1 = \frac{e^{(P_1G - C_1)/\tau}}{e^{(P_1G - C_1)/\tau} + e^{((1 - P_1)G - C_2)/\tau}}.$$

Now we shall describe the asymptotes of this choice probability and resulting behaviour of the system at extreme values of G and τ . These properties have been shown and then tested on a model in (Belavkin 1999).

- i) $\tau \rightarrow 0$ (no noise). In this case the system entirely relies on statistical information and may be too deterministic. We speculate that the behaviour models in some cases the behaviour that has been described by Damasio concerning some of his patients, such as repetitive errors, inability to choose from equal opportunities. The first is due to excessive reliance on the past experience, which may become obsolete in a changing environment because it takes too long to learn new statistics to override the old ones. The second type of behaviour occurs simply due to the possibility of equal expected gains for several rules.

ii) $\tau \rightarrow \infty$ (high noise). It is easy to show that

$$p_1 \rightarrow \frac{1}{2}, \quad \forall P_i, C_i$$

In this case the choice becomes completely random (or irrational) since it does not depend on the past experience at all. Note, that sometimes such behaviour may be useful.

iii) $G \rightarrow 0$ (low motivation). Then

$$p_1 \rightarrow \frac{e^{-C_1}}{e^{-C_1} + e^{-C_2}}.$$

In this case the choice is completely determined by the the costs C of rules and not by the expected probabilities P . The system is trying to put as little efforts into the task as possible and does not “care” about the probability of successful outcome.

iv) $G \rightarrow \infty$ (too high motivation). In this case the resulting p_1 will depend on the value of expected probability P_1 , which can be between 0 and 1. In the extreme cases we shall have:

$$\begin{aligned} p_1 &\rightarrow 1 && \text{for } P_1 = 1 \\ p_1 &\rightarrow 0 && \text{for } P_1 = 0. \end{aligned}$$

This case is opposite to the previous one: the choice does not depend on the costs C , but is purely determined by the expected probabilities P . The system does not pay attention to the effort (time) it spends, and achieves the goal whatever the cost.

Note, that if we increase both G and τ keeping the ratio G/τ constant, then expected probabilities become more important than costs. Indeed, let $\tau \rightarrow \infty$ and $G \rightarrow \infty$. In this case, like in iii), asymptotes of p_1 are determined by the value of P_1 :

$$\begin{aligned} p_1 &\rightarrow \frac{e}{e+1} && \text{for } P_1 = 1 \\ p_1 &\rightarrow \frac{1}{e+1} && \text{for } P_1 = 0. \end{aligned}$$

So, p_1 for $P_1 = 1$ is e times bigger than for $P_1 = 0$ ($e > 1$).

Similarly, for both $G \rightarrow 0$ and $\tau \rightarrow 0$ the costs become more important. It means that the values of G and τ are important and not only their ratio G/τ .

4 On activation theory of emotion

Asymptotic properties of choice probability show that the ratio G/τ determines how much the choice depends on the learned statistical information, while the values of G and τ determine whether the choice is made from costs of probability perspective. Based on this observation we may think of the G/τ ratio as an indicator of confidence

of a problem solver and the values of G and τ representing the activation or the “energy” of a cognitive process, which is called *arousal* in activation theory of emotion. If this is true, then the laws of activation theory of emotion, such as Yerkes-Dodson Inverted-U curve relating arousal to performance, should also apply to ACT-R cognitive models.

4.1 The Yerkes-Dodson experiment

In order to test this idea, a model replicating the classical “dancing mouse” experiment by Yerkes and Dodson (Yerkes & Dodson 1908) was built (Belavkin & Ritter 2000) using ACT-R.

In the original experiment a mouse was placed into a discrimination box with two exits: one marked by a white card and another by a black one. A mouse was trained initially to exit the box through any door, but after two days it was only allowed to exit through the white door and if it did a mistake, it was subjected to a slight electric shock at the black door. The order of the doors was changed randomly and researchers measured the number of wrong choices a mouse made each day until a perfect habit had been formed (i.e., when no errors were produced for three consecutive days).

Such experiments with certain variations were performed by many psychologists studying discrimination learning or perception in the early 20-th century. The remarkable feature of Yerkes and Dodson work was that they looked at the speed of learning with respect to the strength of stimuli. They changed the lighting of the box so that it was easier or harder to discriminate between the white and the black doors, and they increased the strength of the electric shock from weak to medium and strong.

The main result of this experiment was that the best (fastest) learning occurred under the medium stimulus (Figure 3. From (Yerkes & Dodson 1908)). Performance for the strong stimulus was better than for the weak one, but it was worse than medium, especially when visual discrimination was not perfect.

4.2 Task simulation and “dancer” model

The task simulation was implemented using Common Lisp and Garnet graphics library. The user interface consists of three main windows (Figure 4): discrimination box with three rooms and two doors, a control panel with a slider for setting the contrast between the two doors and a control for setting the virtual “voltage”, and the third is a window for displaying the number of errors a mouse is making and other data. The contrast between the doors is related to the G/τ ratio in ACT-R while the voltage control indirectly sets the goal value G .

The dancing mouse is represented by a red arrow head object and its actions are controlled by a cognitive model, implemented with ACT-R. It uses a simple perception-action cycle and two-dimensional world representation model.




Figure 3: Learning curves for weak “W”, medium “M” and strong “S” stimuli under medium visual discrimination. Abscissae represent the average number of errors for four mice produced in ten tests for each series. Ordinates represent the number of series.



Figure 4: User interface of the “dancing” mouse experiment simulation. Left: the discrimination box with two doors. Right: a simple control panel with a slider for setting a contrast between the doors (and the G/τ ratio) and voltage control related to G value.

It has visual sensors sending information about the object right ahead and skin sensors sending the information about the strength of external stimulus. The current and goal states are represented by objects of a special type *self*, that can be located at some point or inside an environment (a room) object. The mouse model uses means-ends analysis concept to achieve a goal by taking actions (turning and moving) to reduce the difference between the current and the goal state.

The model uses ACT-R’s two main learning mechanisms: *probability learning* (statistical learning about the usage of a particular production rule) and *production compilation* mechanism (Anderson & Lebiere 1998) to form new production rules.

The main room the mouse is initially placed in has two exits, and the choice of one of two objects is repre-

sented by a special chunk-type *choice*. When the model is presented with a choice chunk as a goal, it has initially only two actions encoded in two simple production rules like the one below to choose the first object:

IF the goal isa *choice* of *first* or *second*
THEN focus on *first*

and a similar rule for choosing the second object. Here no features of the *first* and *second* objects are used to make the decision.

With no electric shock behind the doors the mouse chooses doors randomly and due to the statistical probability learning it may eventually form a slight preference for one of the doors. When the shock is introduced, on choosing the wrong door (black door) the mouse recalls the last choice it has made and learns to choose another alternative paying now attention to the features of the objects it is choosing. This learning uses the production compilation mechanism and it may add a new production rule like:

IF the goal isa *choice* of *first* or *second*
AND *first* isa *black door*
THEN focus on *second*

So, if next time the mouse is presented with the same choice again, it will already have three production rules. Note, that in the above rule the learning occurs based on one dimension — colour. Another dimension is door position (left or right) and a mouse can learn to choose based only on the position of the door. Two-dimensional learning, when both colours and positions are taken into account, is also possible. The probability of using the colour information in learning depends on the contrast between the doors.

The model was tested under several scenarios. With the contrast control set to a particular value, the G/τ ratio remained constant, while the values of G and τ were controlled by the “voltage” going from small to high values.

In the first experiment only the speed of probability learning was studied since G and τ parameters affect only the conflict resolution and not the production compilation. So, the doors remained in the same order (statistical learning would not work if the order of the doors was random) and information about eventual successes or failures was used to learn the expected probabilities of the two competing rules.

As predicted by the asymptotic properties, probability learning is badly affected by high noise (low G/τ), but the performance increased with the “voltage”. The later is due to the fact that probabilities P become more important at high G values (even for the same G/τ). The curve for number of errors during statistical learning has a shape of exponential decay (Figure 5). The speed of errors decay increases with lower noise temperature τ and higher goal value G . At extremely high values of G and τ and when the ratio G/τ is not very high (significant noise) the




Figure 5: Model performance in the first experiment on the speed of statistical learning. The error curves have exponential decay shape with the fastest decay for medium (“M”) stimulation. See Figure 3 for more description.

performance becomes more varied (higher standard deviation) and the average at least does not become any better.

In the second experiment the model was tested with the production compilation mechanism and a random ordering of the doors. The resulting curves have slightly different shapes with a distinguished “break point” and slope. This is due to the fact that statistical probability learning for the initial two rules does not produce useful results since the doors are changed randomly. The break point corresponds to the moment when the learned new rules, including information about the features of the doors, start being used. Initially, the newly formed production rules have a small strength, but their chance to win the competition in conflict resolution increases every time the rule is relearned. The more often a rule is used the more information about its eventual success or failure is being collected and the more important becomes its probability P . This way a mouse learns to avoid the black door.

Nevertheless, the effect of G and τ on conflict resolution is still important since all the rules satisfying the choice goal are in the conflict set. The curves for a different strength of stimuli are shown on Figure 6.

The higher degradation of performance could be better observed if we assumed that the ratio of G/τ decayed at high “voltage” values. Recall that this ratio can be related to confidence, which obviously did not grow under the fear of a strong electrical shock. In addition, there might be other factors such as memory loss or perception damage affecting the performance under strong stimulation.



Figure 6: Model data for the second experiment with random doors order and production rules learning. The error curves decay faster when the newly learned rules become stronger (a break-point). See Figure 3 for more description.

5 Dynamics of motivation and randomness during problem solving

In this section the probabilistic interpretation of expected gain E will be introduced and ACT-R conflict resolution will be considered as a probability maximisation problem. We shall discuss the dynamics of the G and τ parameters and its analogy with emotions experienced during problem solving.

5.1 Parameters game

The dancing-mouse model showed that the Yerkes-Dodson Inverted-U curve effect can be observed on cognitive models and that we could model different arousal and motivational states using the goal value and noise parameters in ACT-R. But we know that arousal, motivation and confidence may change during the problem solving when we experience emotions such as frustration or joy.

It was proposed earlier in (Belavkin et al. 1999) that the goal value G and noise temperature τ should not remain constant, which was more obvious after some experiments with the Tower of Nottingham model. An attempt of the authors to model 3–4 years old children behaviour by simple further increase of noise did not lead to the desired result: the model could not solve the problem anymore and sometimes it ran for several simulated hours, which obviously was very far from the behaviour of 3–4 years old children. A reasonable question to ask was why the model did not abandon the task?

As was mentioned earlier, the parameter indicating how long it is planned to spend on the task is the goal value G and its default value is 20 minutes. But in ACT-

R by default its value remains constant throughout the task,¹ which means that after unsuccessful two hours the model is still ready to spend another 20 minutes solving a problem. Obviously, this is not reflecting reality. The goal value should be seen as the maximum amount of resources (not necessarily time) that the problem solver is ready to allocate at the current cycle, but it may not, and perhaps, should not remain constant.

5.2 Expected gain and probability

Let us consider a problem solver, that has a set of random decisions $x = \{X_0, \dots, X_n\}$ and it is supposed to find a solution at one of the time moments $t = \{T_0, \dots, G\}$. We use G here for compatibility with ACT-R notation and it is the “dead-line” by which we plan to get the answer.

If a problem is solvable in principle, then it means that initially a problem solver has enough knowledge and means to interact with the task to be able reach the goal state and given the same problem again it should be able to repeat this process. In other words, if a goal is achievable, then it means that there exists at least one decision $X \in x$, such that applying it to the problem at the initial time moment will lead eventually to the goal state at a moment $C \in t$ (again, we use C according to ACT-R notation for the cost).

Let us denote by y some states in the environment or working memory (problem space) and let $Y \in y$ be such a state, that satisfies the goal criteria (in other words, Y is the goal state). We can think of Y as an evaluation of some function at the (X, C) point:

$$Y = f(X, C, \xi_1, \dots, \xi_n),$$

where ξ_i are some random variables (unknown parameters).

Since the function f may be not known and there may be some unknown parameters, we may consider a probability $P(X, C | Y)$ of the decision X and time moment C for the desired outcome Y .² This probability is *a priori*, and $P_{\text{out}}(X, C | Y) = 1$ is *a posteriori* probability if the goal state Y has been achieved. If decisions x and time moments t are conditionally independent, then

$$\begin{aligned} P_{\text{out}}(X, C | Y) &= P_{\text{out}}(X | Y)P_{\text{out}}(C | Y) \\ &= P_{\text{out}}(X | Y) = P_{\text{out}}(C | Y) = 1, \end{aligned}$$

where $P_{\text{out}}(X | Y)$ and $P_{\text{out}}(C | Y)$ are marginal probabilities. Since $G \geq C$ for $C \in \{T_0, \dots, G\}$ we can write the following inequality:

$$P_{\text{out}}(X | Y) \frac{G}{C} - 1 \geq 0.$$

or

$$P_{\text{out}}(X | Y)G - C \geq 0,$$

¹Here we do not consider subgoals and their values.

²Here P is not the expected probability of a rule in ACT-R, because it considers both rules and time moments

which corresponds to the form of equation (1) for the expected gain.

We can see now that G and C in (1) represent the “time” component of the *a priori* probability $P(x, t | Y)$ and choosing a rule by maximum expected gain means choosing by the maximum *a priori* probability of rules and their costs for a particular goal state. The ratio G/C is proportional to the marginal probability $P(C \leq G | Y)$ and if it is not 0, then by increasing G we increase the chance to find the solution.

5.3 Climbing the probability hill

The values of the *a priori* probabilities P for different rules and time moments can be represented on a 3D graph. For a better illustration, we may consider a probability density surface $\varphi_Y(x, t)$ on continuous decisions-time plane for a goal state Y (Figure 7). The points (x, t) on this surface correspond to expected probabilities P and costs C of production rules in ACT-R. This information is learned and updated during problem solving. Finding the maximum on this probability surface is a hill-climbing task with G determining the angle, or direction of the search. A solution for the direction (value of G) is given by the maximum-gradient method, and ideally it should change towards the maximum incline. But the maximum gradient method does not necessarily give a unique solution and the initial direction.




Figure 7: 3D representation of hill-climbing analogy. The surface φ is the probability density of decisions x and time moments t for a goal state Y . The line (P, C) represents current information about P s and C s of the production rules. The line $t = Gx$ represents the direction or area of search for the given G value. Three lines represent different cases: e^0 when $G = \text{constant}$ throughout the task; e^+ when G varies towards the maximum incline; e^- when G varies from the incline.

It is not hard to notice that low G corresponds to breadth-first search. Indeed, according to asymptotic properties, low G gives most priority to rules with low costs (less

time needed), and it means that more production rules can be tried. On the other hand, high G corresponds to depth-first search since it gives priority to rules with high expected probabilities P not considering their cost very much (problem solver may spend more time executing one rule with higher cost). So, problem solving with low motivation corresponds to breadth-first search, while high motivation to depth-first search. A search method combining these two strategies is known as best-first search (from breadth to depth), and it suggests that G (motivation) should gradually increase during the problem solving approaching a goal.

5.4 Annealing analogy

On the contrary, noise temperature τ should be high in the initial state of problem solving making the process more random. This corresponds well to the situation, when subjects start solving a task, such as Tower of Nottingham, when they are just playing with blocks and trying some simple constructions. When more information becomes known about the task, subjects become more confident, which corresponds to higher G/τ ratio and lower noise.

Now, looking at the Boltzmann equation (2), we may notice that such heuristics for controlling G and τ is exactly the same as optimisation by simulated annealing (Kirkpatrick, Gelatt & Vecchi 1983) (with E playing a role of negative energy).

5.5 Emotion and heuristics

This observation suggests that an emotional problem solver in general follows a very powerful optimisation methods:

- Positive emotions, experienced on successes during problem solving, are accompanied by increase of the motivation (goal value G) and confidence (G/τ ratio). This process corresponds to cooling the system in the simulated annealing and a goal state corresponds to crystallisation.
- Negative emotions correspond to a decrease in G and G/τ (heating the system up). Negative emotions occurring during problem solving can play a positive role in overcoming possible problems. In a “hill-climbing” illustration these problems are known as *local maximum*, *plateau* and *ridge*. The strategies used to overcome these problems are the change of direction of the search (G decrease) and random jump (noise temperature τ increase). In simulated annealing it corresponds to melting the system from a glass state.

The amount of overall activation before experiencing success or failure may determine the strength of the experienced emotion. For example, frustration corresponds to negative emotion with low activation, while anxiety to negative with high activation.

The relation of positive and negative emotions with high or low value of G/τ ratio correlates very well with another observation that subjects tend to overestimate a success and underestimate a failure in a good mood, while being more sceptical in a bad mood (Johnson & Tversky 1983) (Nygren, Isen, Taylor & Dulin 1996). Indeed, at high values of G and low τ the expected gain (1) depends entirely on the statistically learned expected probability P , which may not necessarily reflect the real situation. The reverse situation occurs at low G and high τ when the learned choice does not depend on P , although its value may perhaps be 1.

6 Conclusion and open questions

In this work we tried to apply modern cognitive science and modelling methods to the subject relating intelligence and emotion. The subject is highly speculative and there are still many questions with no answers. For example, how to evaluate a success or failure during problem solving and how to implement it in cognitive architecture to make the emotional changes in information processing automatic? If motivation control plays the role of resource management, what should be the uniform measure for the resources of an artificial intelligence with emotions?

These questions hopefully will be answered in the foreseeable future, but now let us try to summarise the results of this work. Three main conclusion can be:

1. It is possible, and we have shown how, to model a range of psychological phenomena, related to emotion using cognitive architectures such as ACT-R.
2. It has been shown that emotion and affect should be taken into account by cognitive modellers since the motivational and emotional changes during problem solving may produce noticeable effects in performance.
3. It has been shown that in general emotion makes a positive contribution to problem solving, since it implements powerful heuristic methods already known in AI and mathematics, and hence it is important for intelligence.

Acknowledgements

I would like to thank David Elliman, David Wood, Frank Ritter and Brian Logan for helpful discussions and supervision of this work. This work is sponsored by EPSRC Credit and ORS Award Scheme.

References

- Anderson, J. R. (1993), *Rules of the Mind*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Anderson, J. R. & Lebiere, C. (1998), *The Atomic Component of Thought*, Lawrence Erlbaum Associates, Mahwah, NJ, London.
- Baxter, G. D. & Ritter, F. E. (1996), Designing abstract visual perceptual and motor action capabilities for use by cognitive models, Technical Report 36, ESRC CREDIT, Psychology, U. of Nottingham.
- Belavkin, R. V. (1999), Annual report: studying the role of emotions in problem solving with cognitive architectures, Technical report, School of Computer Science and Information Technology, University of Nottingham.
- Belavkin, R. V. & Ritter, F. E. (2000), Adding a theory of motivation to ACT-R, in 'Seventh annual ACT-R workshop proceedings', Carnegie Mellon University, Department of Psychology, Pittsburg, PA 15213, pp. 133–139.
- Belavkin, R. V., Ritter, F. E. & Elliman, D. G. (1999), Towards including simple emotions in a cognitive architecture in order to fit children's behaviour better, in N. Mahwah, ed., 'Proceedings of the 1999 Conference of the Cognitive Science Society', Lawrence Erlbaum, p. 784.
- Damasio, A. R. (1994), *Descartes' Error: Emotion, Reason, and the Human Brain*, Gosset/Putnam Press, New York, NY.
- Friedman, M. P., Burke, C. J., Cole, M., Keller, L., Millward, R. B. & Estes, W. K. (1964), Two-choice behaviour under extended training with shifting probabilities of reinforcement, in R. C. Atkinson, ed., 'Studies in mathematical psychology', Stanford University Press, Stanford, CA, pp. 250–316.
- Goleman, D. (1995), *Emotional Intelligence*, Bantam Books, New York.
- Johnson, E. & Tversky, A. (1983), 'Affect, generalization, and the perception of risk', *Journal of Personality and Social Psychology* **45**, 20–31.
- Jones, G., Ritter, F. E. & Wood, D. J. (2000), 'Using a cognitive architecture to examine what develops', *Psychological Science* **11**(2), 1–8.
- Kirkpatrick, S., Gelatt, C. D. & Vecchi, J. M. P. (1983), 'Optimization by simulated annealing', *Science* **220**(4598), 671–680.
- Lebiere, C. & Anderson, J. R. (1993), A connectionist implementation of the ACT-R production system, in 'Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society', Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 635–640.
- LeDoux, J. E. (1996), *The Emotional Brain*, Simon & Schuster, New York.
- Myers, J. L., Fort, J. G., Katz, L. & Suydam, M. M. (1963), 'Differential monetary gains and losses and event probability in a two-choice situation', *Journal of Experimental Psychology* **77**, 453–359.
- Newell, A. (1990), *Unified Theories of Cognition*, Harvard University Press, Cambridge, Massachusetts.
- Nygren, T. E., Isen, A. M., Taylor, P. J. & Dulin, J. (1996), 'The influence of positive affect on the decision rule in risk situations', *Organizational Behavior and Human Decision Processes* **66**, 59–72.
- Ortony, A., Clore, G. L. & Collins, A. (1988), *The Cognitive Structure of Emotions*, Cambridge University Press, Cambridge, MA.
- Picard, R. W. (1997), *Affective Computing*, MIT Press, Cambridge, Massachusetts. London, England.
- Plutchik, R. (1994), *The Psychology and Biology of Emotion*, 1 edn, HarperCollins College Publishers, New York.
- Roseman, I. J., Antoniou, A. A. & Jose, P. E. (1996), 'Appraisal determinants of emotions: Constructing a more accurate and comprehensive theory', *Cognition and Emotion* **10**(3), 241–277.
- Slooman, A. (1999), 'Review of affective computing', *AI Magazine* pp. 127–133.
- Yerkes, R. M. & Dodson, J. D. (1908), 'The relation of strength of stimulus to rapidity of habit formation', *Journal of Comparative and Neurology and Psychology* **18**, 459–482.