In C. Johnson (Ed.), *Proceedings of the AISB'04 Symposium on Emotion, Cognition and Affective Computing*
ISBN 1-902956-36-6, (pp. 1–8). Leeds, UK, 2004.

On Relation between Emotion and Entropy

Roman V. Belavkin
*School of Computing Science, Middlesex University
London NW4 4BT, UK
R.Belavkin@mdx.ac.uk

Abstract

The ways of modelling some of the most profound effects of emotion and arousal on cognition are discussed. Entropy reduction is used to measure quantitatively the learning speed in a cognitive model under different parameters' conditions. It is noticed that some settings facilitate the learning in particular stages of problem solving more than others. The entropy feedback is used to control these parameters and strategy, which in turn improves greatly the learning in the model as well as the model match with the data. This result may explain the reasons behind some of the neurobiological changes, associated with emotion and its control of the decision making strategy and behaviour.

1 Introduction

It is popular to believe now that emotion is an important (if not essential) component of intelligence (Salovey and Mayer, 1990). This is, however, hard to prove unless some quantitative methods are introduced that will allow us to evaluate such claims in an experiment. An example of such an experiment could be a competition between several agents, with architectures incorporating various theories of emotion and cognition. In practice, however, the results of such an experiment would be very hard to interpret because of the great number of components (e.g. perception, memory, planning, action, etc) involved in the agents' architectures.

The research described in this paper pursues a different approach by studying the effects of emotion on decision making and learning. Using entropy reduction as a quantitative measure of learning allows for a better analysis and comparison of the results from different experiments.

The ability to learn is one of the most important features of intelligent systems. While leaving to philosophers the question of what is the purpose of learning, let us assume that this process is beneficial to intelligent systems, and the faster and more effectively it occurs the better. From information theory point of view learning is equivalent to reducing the uncertainty (entropy) about the environment and the system itself within this environment. Many areas of artificial intelligence have already successfully employed the mathematical apparatus of information theory, which advanced greatly the neural networks learning algorithms, search methods and case–based reasoning systems. Recently, the notions of information and entropy have been applied to analyse and control cognitive models (Belavkin and Ritter, 2003). In particular, it became possible for the models implemented in hybrid cognitive architectures, such as ACT–R (Anderson and Lebiere, 1998), which mixes the high level symbolic processing with the low level subsymbolic computations accounting for fuzzy or probabilistic properties of cognition.

The comparison of model results with data (e.g. from human subjects or animals) is one of the most important aspects of the cognitive modelling research. A cognitive model of a classical animal learning experiment will be used in this study to evaluate theoretical predictions.

In the next Section, the most general effects of basic emotions and arousal on behaviour will be discussed and grounded in the relevant literature. The ambiguity of the term emotion will be avoided by replacing it with the principle components of emotions.

The notion of entropy and its application to cognitive models will be discussed in Section 3. This section will repeat some of the previous work (Belavkin and Ritter, 2003). Section 4 will highlight how speed of learning in the model varies as a function of some parameters in the architecture. These parameters (namely the noise variance and goal value used in decision making mechanism) have been used before to simulate different levels of motivation and arousal (Lovett and Anderson, 1996; Anderson and Lebiere, 1998; Belavkin, 2001). The entropy reduction will be used to measure the speed of learning in the model.

Section 5 will discuss the idea of using the entropy of success as a feedback parameter to control the decision making mechanism of the architecture. It will be shown how the entropy evaluating model's own performance moderates the choice strategy and controls the behaviour making it more adaptable. In addition, the model's match with the data improves, which supports the idea that a similar strategy control takes place in subjects. Some more speculative ideas about the role of emotion in evaluating the entropy and controlling the behaviour will be discussed in the end of the paper.

2 The Principle Components of Emotions

The important role of emotion in cognition has been extensively discussed in the literature, particularly over the last two decades (Salovey and Mayer, 1990; Damasio, 1994; LeDoux, 1996). Despite the great interest in the subject of emotion across several disciplines of science, there is still a lack of understanding and clear definition of what emotion actually is. Psychologists and philosophers still cannot agree on some of the fundamental points in the subject, such as what comes first: Feelings or thought? (Schachter and Singer, 1962; Zajonc, 1980).

This ambiguity is multiplied when one attempts to integrate emotion into a unified theory of cognition, and into its computational implementations, such as ACT–R (Anderson and Lebiere, 1998) or SOAR (Newell, 1990). The need to include emotion into cognitive models, however, is rarely disputed (Simon, 1967). With the existence of many computational models of affect (see Hudlicka and Fellous (1996) for a review) and even a greater number of different emotions (Lambie and Marcel, 2002), the problem seems to be intractable. However, the dimensionality can be reduced if we concentrate our research on measurable and the more consistent features of the phenomena, or what we shall call the *principle components of emotions*.

Probably the most common measure of various emotional experiences is *valence* indicating whether an emotion is positive or negative. Cannon (1929) argued that all emotions can be classified into 'fight or flight', which is probably not far from the truth. Another important measure is *arousal*, or the intensity of emotional experience. Arousal is a broad term covering a variety of phenomena, but generally it is associated with different levels of activation of the autonomic nervous system (ANS), and it can be influenced by external or internal stimulation including emotion (Humphreys and Revelle, 1984). As has been shown by Russell (1983, 1989), valence and arousal are the two most common dimensions in classifications of emotions, and they are included in many other classifications (Plutchik, 1994).

Both valence and arousal are measurable and even predictable. Indeed, negative emotions occur when we experience a failure in achieving a particular goal. On the contrary, a success is accompanied by positive emotions. Arousal can be either measured directly in subjects (e.g. using galvanic skin response), or predicted based on the strength of the stimuli (e.g. reward or penalty). Therefore, in this paper, when discussing the role of emotion in cognition, we shall concentrate on the effects of arousal and valence, and we shall not consider other aspects of the phenomenon, such as particular emotions or their role in social interaction and so on.

On an individual level, emotion is known to play a role in different aspects of cognition, such as perception, memory, action and learning (LeDoux, 1996). There is quite a lot of experimental evidence suggesting the relation between arousal and cognitive performance. For example, the studies of the inverted–U effect showed the relation between arousal and the speed of learning (Yerkes and Dodson, 1908; Mandler and Sarason, 1952; Matthews, 1985). Another series of experiments showed how the expectation of positive or negative outcomes may change the decision making strategy (Tversky and Kahneman, 1981; Johnson and Tversky, 1983). Below is the summary of some effects of valence and arousal that can be useful in designing a cognitive model:

- Positive valence is associated with success, choice involving gains and risk aversive behaviour. Negative valence is associated with failure, choice involving losses, and the behaviour is usually more risk taking (Tversky and Kahneman, 1981; Johnson and Tversky, 1983).

- Low arousal is associated with low level of stimulation or motivation, actions requiring less efforts are more likely. High arousal is associated with high level of stimulation or motivation, actions involving more efforts are more probable (Humphreys and Revelle, 1984).

It has been suggested before (and will be discussed in Section 4 of this paper) how to achieve the above types of behaviour in cognitive models using parameters manipulation (Belavkin, 2001). The speed of learning in the model under these parameters settings will be measured by means of entropy reduction. In the next section, we discuss some definitions of entropy and an example of calculating it in a cognitive model.

3 Information and Learning

Learning is one of the most important characteristics of intelligence. It allows a subject or a system to improve the performance in certain tasks or class of problems. The most obvious measure of such an improvement is an increase of success rate, or equivalently a reduction of failures (errors). Ultimately, learning reduces the uncertainty of the outcome with the success being a more probable. Thus, entropy reduction could be a convenient measure of learning. However, in practice it is impossible to measure directly in subjects the parameters necessary for entropy computations (e.g. synaptic weights), and traditionally learning is judged based on external observations (i.e. the reduction of errors such as shown on Figure 1).

Unlike the brains of subjects, however, cognitive architectures allow for a relatively easy access to all the internal variables. This opened a possibility to measure the learning in cognitive models directly by calculating the entropy change or information (Belavkin and Ritter, 2003). The advantage of using entropy is that it provides a compact display of the internal changes in a model as a result of learning, which may not always have external manifestations. In this section, the use of entropy to describe learning in intelligent systems will be described and shown on an example of a cognitive model.

3.1 Entropy and surprise

In the most general case, entropy H is a monotonic function describing the complexity (or uncertainty) of a system, such as $H = \ln M$, where M is the number of states a system can be in. This canonical definition assumes no information about the probabilities of individual states. If, however, we know the probabilities $P(\xi)$ of different (random) states ξ, then the entropy can be calculated as:

$$H(\xi) = -E\{\ln P(\xi)\} = -\sum_{\xi} P(\xi) \ln P(\xi)\,, \qquad (1)$$

where $E\{\cdot\}$ denotes the expected value operator. If all states ξ are equally probable, then entropy (1) equals $\ln M$, and it corresponds to the maximum value of H for given M. Thus, the uncertainty can be reduced if by means of Bayesian estimation we find out which states have greater likelihood. Shannon (1948) defined information as the difference between entropy before and after an observation of some event y:

$$I(x, y) = H(x) - H(x \mid y)$$

Here, x denotes some variable, the information about which is received indirectly through observation of y.

Interestingly, information and entropy have been used before to explain one basic emotion — surprise. Indeed, the lower is the probability P of event ξ, the greater is the amount of information $-\ln P(\xi)$ received when this event happens (i.e. the greater is the surprise). This early observation points to the possibility that our nervous system and body reacts to the amount of information received, and the feedback seems to be proportional to this amount. Note, however, that surprise can be positive as well as negative, and the reaction can be different in each case. In this paper, we shall look more carefully into the nature of such a feedback, and investigate using a cognitive model whether this feedback is beneficial for an intelligent system (i.e. helps in learning and adaptation).

3.2 Uncertainty of success

It is quite difficult to estimate the entropy of a large system with many states (e.g. a cognitive model). However, for an intelligent system it is possible to look at the problem from a different perspective: The uncertainty of whether it achieves the goal or not (Belavkin and Ritter, 2003). The *entropy of success* has been defined as

$$H_{\mathrm{SF}} = -\left[P(\mathrm{F}) \ln P(\mathrm{F}) + P(\mathrm{S}) \ln P(\mathrm{S})\right]\,, \qquad (2)$$

where $P(\mathrm{S})$ is the probability of success in achieving the goal, and $P(\mathrm{F})$ is the probability of failure. Note that $P(\mathrm{F}) = 1 - P(\mathrm{S})$. If a system (e.g. a cognitive model) has to choose from a set of n alternative decisions to achieve the goal, then the probability of success is:

$$P(\mathrm{S}) = \sum_{i=1}^{n} P(\mathrm{S}, i) = \sum_{i=1}^{n} P(\mathrm{S} \mid i) P(i)\,, \qquad (3)$$

where $P(\mathrm{S}, i)$ is the joint probability of successful outcome and ith decision, $P(\mathrm{S} \mid i)$ is the conditional probability of success given that ith decision has been made, and $P(i)$ is the probability of ith decision. Thus, to calculate the entropy of success H_{SF}, one should estimate probabilities $P(\mathrm{S} \mid i)$ and $P(i)$, which depend on specific architectural implementation (i.e. SOAR, ACT–R, neural networks, etc).

Conditional probabilities $P(\mathrm{S} \mid i)$ represent the prior knowledge about the likelihood to achieve a success, if certain decisions (and associated actions) are taken. Note that a problem solver may not be aware of or not considering some decisions initially. However, the number of decisions n to choose from may increase with time as the result of learning. Probability $P(i)$ depends on the way the decision making (e.g. rule selection algorithm) is implemented. Thus, $P(i)$ is more related to the architecture rather than the knowledge of a system. As an example, let us consider the ACT–R cognitive architecture (Anderson and Lebiere, 1998).

3.3 Computation of entropy in ACT–R

ACT–R (Anderson and Lebiere, 1998) is a general purpose hybrid cognitive architecture for developing cognitive models that can vary from simple reaction tasks to simulations of pilots navigating airplanes and operators of airtraffic control systems. ACT–R follows the approach of *unified theories of cognition* (Newell, 1990), in which several theories about different aspects of cognition are used in a single simulation system. Today, ACT–R has emerged as the architecture of choice for many cognitive modelling problems.

In ACT–R, decisions are encoded in a form of production rules, and during the model run the number of successes and failures of each rule is recorded by the architecture. This information is used to estimate empirically the probabilities $P(\mathrm{S} \mid i)$ of success for ith rule:

$$P(\mathrm{S} \mid i) \approx P_i = \frac{\mathrm{Successes}_i}{\mathrm{Successes}_i + \mathrm{Failures}_i}\,. \qquad (4)$$

Here P_i is statistics of ith rule. In addition, ACT–R records the efforts (i.e. time) spent after executing the rule and actually achieving the goal (or failing). This information is used to calculate the average cost C_i of ith rule. Parameters P_i and C_i represent subsymbolic information about the decisions, and can be learned statistically. On symbolic level, a model can learn new rules as well as new facts used by these rules.

When several alternative rules are available that match the current working memory state (i.e. the current goal, perception, retrieved facts), then one rule has to be selected using the conflict resolution mechanism. In ACT–R, this is done by maximising the expected utility of rules in the conflict set: $i = \arg\max U_i$, where

$$U_i = P_i G - C_i + \xi(\sigma^2)\ . \tag{5}$$

The above equation has allowed ACT–R to model successfully some important properties of human (and animals) decision making: Probability matching (use of P_i in utility); The effect of a payoff value (G represents the goal value); Stochasticity (the utility is corrupted by zero–mean noise of variance σ^2) (Anderson and Lebiere, 1998).

Although there are other mechanisms in ACT–R, such as chunks (facts) retrieval, that may affect rules' selection, the probability $P(i)$ that ith rule will be chosen can be approximated by Boltzmann equation as:

$$P(i) \approx \frac{e^{\bar{U}_i/\tau}}{\sum_{j=1}^{n} e^{\bar{U}_j/\tau}}, \tag{6}$$

where $\bar{U}_i$ is the utility not corrupted by the noise, and $\tau = \sqrt{6}\sigma/\pi$ is called the *noise temperature*. Now, using approximations (6) and (4), one can calculate the success probability (3) and entropy of success (2).

3.4 A model example

The reduction of entropy of success has been used to analyse the learning in an ACT–R model of the Yerkes and Dodson (1908) experiment (Belavkin, 2003). In this classical experiment, mice were trained over several days to escape a discrimination chamber (a box with two doors) from one particular door, and the number of errors was measured for every day. Figure 1 shows an example of the learning curve representing the number of errors produced by the model in this task during 10 tests per each simulated day. The learning curve, however, does not provide a very detailed picture of what is learned and when.

The performance of the model improves because it learns new production rules, and then by trying these rules the model updates their statistics (P_i and C_i) and uses the most efficient and effective ones. Figure 2 shows the traces of probabilities P_i of production rules relevant to the problem goal in the same experiment. One can see that as new rules and statistics are learned after Day 1, the number of errors decreases (see Figure 1). However, the model produces more errors during Days 5, 6 and 7, which means that the model did not have sufficient knowledge, and the errors forced the model to learn more rules. The model learned new rule during Day 5, but the trace of its statistics indicates that the rule was not very helpful (probability of success quickly decreased to $P_i \approx .5$). The new rules learned on Day 7 turned out to be more successful, and the model did not produce any errors after Day 8. One can see that the probability trace reveals much

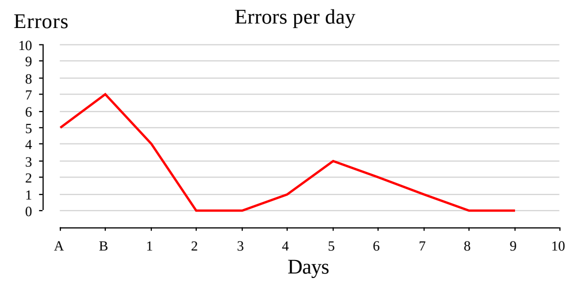

Figure 1: Error curve produced by the model in one experiment.

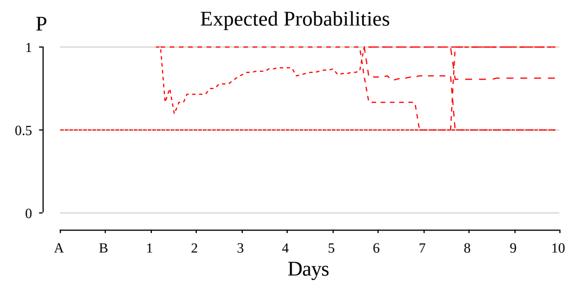

Figure 2: Dynamics of probabilities of rules matching the problem goal. The number of curves increases as new rules are being learned.

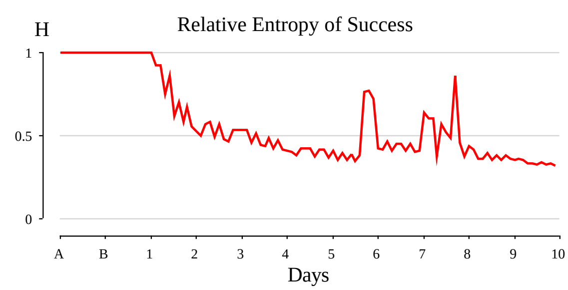

Figure 3: Relative entropy of success of the choice rules. Entropy increases on errors (see Figure 1) and when new rules are learned.

more about the learning in the model than the number of errors.

Figure 3 shows the dynamics of relative entropy of success (relative to the maximum entropy $\ln 2$), calculated using equations (4) and (6) over the probabilities of rules shown on Figure 2. The entropy clearly decays over time indicating the amount of information gained by the model. Also, the entropy increases when the model produces errors, which confirms the idea that entropy of success predicts how certain is the outcome. However, one may notice that the entropy increases most dramatically when new rules are learned (i.e. Days 5 and 7). This can be explained as follows. When new rules are created, the number n of decisions increases, thus making the system more complex (recall that entropy is a function of the number of states). Moreover, the probabilities $P(\mathrm{S} \mid i)$ of the new rules initially have default prior estimates (e.g. .5), and they can only be updated statistically after their application. If the new rules improve the performance, then the entropy of success reduces again (see Day 8, Figure 3).

This example illustrates how entropy change or information can be used as a quantitative measure of learning in a cognitive model. In the next section, the entropy will help analyse how the speed of learning in the model varies as a function of parameters settings in the ACT–R architecture.

4 Variable speed of learning

In ACT–R, the choice of decisions does not depend only on the statistical information about the rules (i.e. estimates of probabilities). Indeed, choice probability (6) depends also on two global parameters in the architecture: The amount of noise (noise variance σ^2 parameter) and the goal value G used in the utility equation (5). Asymptotic analysis of choice probability as a function of σ^2 and G has suggested how different levels of arousal and valences can be simulated in an ACT–R model (Belavkin, 2001):

- At a low noise variance σ^2, the choice is more rational and driven by utility maximisation. Thus, it can be well suited for simulation of the risk aversive behaviour typical for choice with positive expectations (Tversky and Kahneman, 1981; Johnson and Tversky, 1983).

- On the contrary, high noise variance leads to a risk taking, irrational choice, which is less defined by utility maximisation. According to Tversky and Kahneman (1981), this is characteristic of choice with high expectation of a negative outcome.

- At a low goal value G, the costs C_i make more significant contribution to the utility (5). Thus, decisions with higher costs are less likely to be chosen. This is suitable for simulating a low arousal state.

- On the contrary, high goal value G is better for simulating a high arousal level, because under these conditions the model is more likely to take costly decisions.

Let us measure how the speed of learning in the model changes under different conditions. We shall use the entropy reduction as a measuring tool. However, because one of the parameters to be changed is noise variance, it is necessary to make the calculation of entropy independent of these changes. This means substituting the choice probability (6), which depends on τ (noise temperature), by a different probability. For example, we can assume that the choice of a rule is completely random: $P(i) = \frac{1}{n}$, where n is the number of rules (decisions). In this case, probability of a success $P(\mathrm{S})$ can be calculated as

$$P(\mathrm{S}) = \frac{1}{n} \sum_{i=1}^{n} P_i \,. \qquad (7)$$

The entropy associated with this probability (calculated similarly by eq. 2) can be used to estimate the knowledge accumulated in the system in the form of empirical probabilities P_i, because it is independent of the way the decisions are made. We refer to this entropy as the *entropy of knowledge* H_{k}.

The experiments showed that H_{k} decays differently under different noise variance settings. It turns out that although noise hinders the performance of the model, at the same time it may help learn faster. Figure 4 illustrates the probability learning in the model for two noise settings: The left plot shows traces of probabilities with low noise ($\tau = 1\%$ of goal value G), and the right plot for high noise settings ($\tau = 20\%$).[1] One can see that at a higher noise settings (top right), probabilities of rules were updated much more often than at a lower noise (top left). Therefore, the model on the right has better estimates of probabilities. Also, the new and probably more successful rules have been learned earlier in the case of high noise.

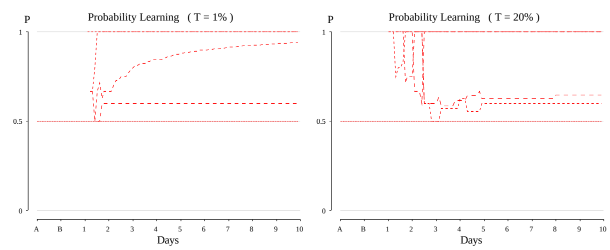

Figure 4: Probability learning under a low noise (left) and a high noise conditions (right).

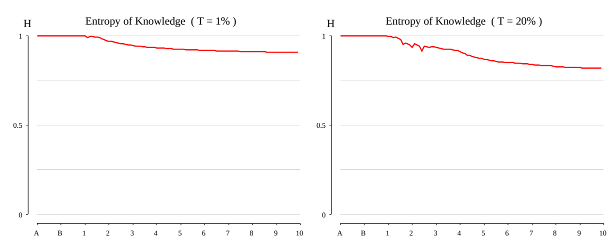

Figure 5: Dynamics of entropy under a low noise (left) and a high noise condition (right).

The corresponding traces of entropies H_{k} are shown on Figure 5. One can see that by day 10 the entropy on the right plot had decayed significantly more than on the left plot. Thus, by day 10 the model with a greater noise gained more information than the model with less noise.

[1] Here noise temperature is calculated as a proportion of the goal value: $\frac{1}{G}\tau \cdot 100\%$.

These results confirm the idea that exploratory behaviour, triggered by an noise increase in ACT–R, facilitates learning in the model.

In the next section, the question of adaptation of behaviour and dynamic control over the parameters in the architecture will be discussed.

5 Entropy feedback and adaptation

The analysis of H_{k} reduction for different noise settings suggested that an intelligent system could benefit from dynamic control over the noise variance. Indeed,

1. At the beginning of solving a problem, exploratory behaviour (high noise) would help gain the information about the task or the environment more quickly.

2. After the important knowledge has been acquired, the choice should concentrate on more successful decisions, which is achieved by the reduction of noise. This should improve the performance.

3. If the environment changes and the number of errors suddenly increases, then a noise increase can speed–up the learning and adaptation of behaviour.

Note that the dynamics of the noise variance, described above, corresponds to the dynamics of entropy in the model (e.g. Figure 3). A simple way to control the noise variance by the entropy parameter has been proposed recently (Belavkin, 2003). More specifically, noise temperature τ was modified in time as:

$$\tau(t) = \tau_0 H_{\mathrm{SF}}(t)\,, \tag{8}$$

where t is time, and $\tau_0 = \tau(0)$ is the initial value of the noise. One can view the noise here as a compensation for the 'missing information', and the otherwise rational, utility–based choice behaviour is corrupted proportionally to the uncertainty.

As predicted, the model with dynamic noise converges faster to a successful behaviour (no errors), and adapts better to changes. What is even more interesting, is that the model fit to the data has improved as well: In one experiment, R^2 increased from $.77$ to $.86$ and the root mean square (RMS) error reduced from 13.2% to 8.8%. Figure 6 shows the learning curves from the static noise model (top) and dynamic noise model (bottom) compared against the data from Yerkes and Dodson (1908). There has been a similar improvement across several data sets Belavkin and Ritter (2003).

The dynamics of noise variance, controlled by the entropy feedback, implements one well–studied heuristics. Indeed, by looking at the Boltzmann equation (6), one can notice that the decrease of noise temperature τ is similar to the optimisation by simulated annealing (Kirkpatrick et al., 1983).

Furthermore, noise variance is not the only parameter in the ACT–R conflict resolution that can optimise the

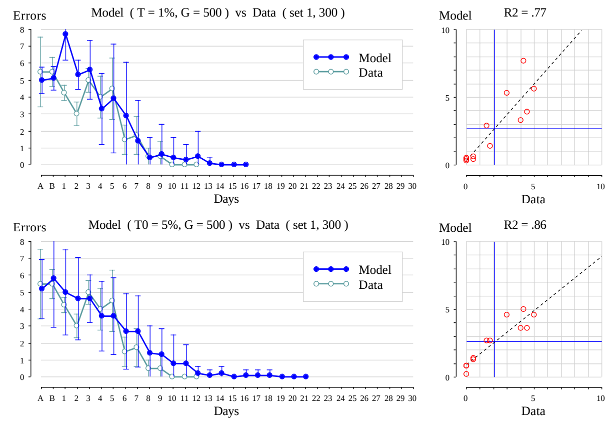

Figure 6: Static noise model (top) and dynamic noise model (bottom) compared with the data (Yerkes and Dodson, 1908). The dynamic model achieves the better match.

learning process. It was shown that goal value G controls the type of the search (Belavkin, 2001): Low G implements the breadth–first search, while high G corresponds to the depth–first search strategy. A search method combining these two strategies is known as the best–first search (from breadth to depth). Thus, gradual increase of G during problem solving can implement the best–first search method.

One can see that the suggested dynamical control of the decision making parameters in the architecture implements some well–known optimisation heuristics, and, therefore, should improve the overall problem solving performance.

6 Discussion

It has been shown in the previous section how dynamic control over two parameters in the ACT–R cognitive architecture improves the learning and adaptive capabilities of the model. In particular, entropy of success has been used as a feedback parameter to control the choice strategy. In addition, this control has improved the match between the model and data. On the other hand, the same parameters have been used to simulate the effects of the principle components of emotions (valence and arousal). Therefore, the dynamic changes of the parameters during problem solving may correspond to the changes in the behaviour due to experiencing emotions of positive or negative valence and the resulting changes of the arousal level. This idea is supported by a number of works in neuroscience and artificial neural networks.

Indeed, in neural networks, the effect of noise can be simulated by changing the bias (or activation threshold) of neurons (Hinton and Sejnowski, 1986). Some neurotransmitters in the brain have a similar effect, and there are areas of the brain (e.g. amygdala) that have connections with the areas of neocortex believed to be responsible for decision–making (LeDoux, 1996).

The role of such interactions have been discussed in the reinforcement learning literature (Sutton and Barto, 1981; Barto, 1985). However, one of the unknown variables there is the amount of reinforcement (e.g. the noise temperature). It has been shown how the entropy of success may help optimise this parameter. Interestingly, entropy and noise temperature have been used for control in the work on analogy by Hofstadter and Marshall (1993).

Today, the idea that emotion plays an important role in controlling and regulating the decision making and actions aspects of cognition is shared by many researchers (Bartl and Dörner, 1998; Sloman, 2001). The results, discussed in this paper, illustrated how the learning in an intelligent system can be improved by using the entropy of success of the system to moderate and control its own behaviour. These observations suggest that appreciation of the system's own performance (entropy of success) and regulating the decision making strategy may indeed be one of the main functions of the emotional system in the brain. Including such an information theoretic feedback mechanism into the design of cognitive models, agent architectures or robots will not only improve their performance, but also will extend our knowledge about the mind and emotion within its context.

References

J. R. Anderson and C. Lebiere. *The Atomic Components of Thought*. Lawrence Erlbaum Associates, Mahwah, NJ, 1998.

Christina Bartl and Dietrich Dörner. PSI: A theory of the integration of cognition, emotion and motivation. In Frank E. Ritter and Richard M. Young, editors, *Proceedings of the Second European Conference on Cognitive Modeling*, pages 66–73, Nottingham, UK, April 1998. Nottingham University Press.

A. G. Barto. Learning by statistical cooperation of self–interested neuron–like computing elements. *Human Neurology*, 4:229–256, 1985.

Roman V. Belavkin. The role of emotion in problem solving. In Colin Johnson, editor, *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing*, pages 49–57, Heslington, York, England, March 2001. AISB.

Roman V. Belavkin. *On Emotion, Learning and Uncertainty: A Cognitive Modelling Approach*. PhD Thesis, The University of Nottingham, Nottingham, NG8 1BB, United Kingdom, March 2003.

Roman V. Belavkin and Frank E. Ritter. The use of entropy for analysis and control of cognitive models. In Frank Detje, Dietrich Dörner, and Harald Schaub, editors, *Proceedings of the Fifth International Conference on Cognitive Modelling*, pages 21–26, Bamberg, Germany, April 2003. Universitäts–Verlag Bamberg.

W. B. Cannon. *Bodily changes in pain, hunger, fear and rage*. Appleton, New York, 1929.

A. R. Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain*. Gosset/Putnam Press, New York, NY, 1994.

G. E. Hinton and T. J. Sejnowski. Learning and relearning in boltzmann machines. In *Parallel Distributed Processing*, pages 282–317. MIT Press, Cambridge, Mass., 1986.

D. R. Hofstadter and J. B. D. Marshall. A self–watching cognitive architecture of high–level perception and analogy–making. Technical Report 100, Indiana University Center for Research on Concepts and Cognition, 1993.

Eva Hudlicka and Jean-Marc Fellous. Review of computational models of emotion. Technical Report 9612, Psychometrix, Arlington, MA, April 1996.

Michael S. Humphreys and William Revelle. Personality, motivation, and performance: A theory of the relationship between individual differences and information processing. *Psychological Review*, 91(2):153–184, April 1984.

E. Johnson and A. Tversky. Affect, generalization, and the perception of risk. *Journal of Personality and Social Psychology*, 45:20–31, 1983.

S. Kirkpatrick, C. D. Gelatt, and Jr. M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598): 671–680, May 1983.

John A. Lambie and Anthony J. Marcel. Consciousness and the varieties of emotion experience: A theoretical framework. *Psychological Review*, 109(2):219–259, 2002.

Joseph E. LeDoux. *The emotional brain*. Simon & Schuster, New York, 1996.

Marsha C. Lovett and John A. Anderson. History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology*, 31:168–217, 1996.

G. Mandler and S. B. Sarason. A study of anxiety and learning. *Journal of Abnormal and Social Psychology*, 47:166–173, 1952.

G. Matthews. The effects of extraversion and arousal on intelligence test performance. *British Journal of Psychology*, 76:479–493, 1985.

Allen Newell. *Unified theories of cognition*. Harvard University Press, Cambridge, Massachusetts, 1990.

R. Plutchik. *The psychology and biology of emotion*. HarperCollins College Publishers, New York, 1st edition, 1994.

J. A. Russell. Two pan–cultural dimensions of emotion words. *Journal of Personality and Social Psychology*, 45:1281–1288, 1983.

J. A. Russell. Measures of emotion. In *The Measurement of Emotions*, volume 4. Academic Press, New York, 1989.

P. Salovey and J. D. Mayer. Emotional intelligence. *Cognition and Personality*, 9(3):185–211, 1990.

S. Schachter and J. E. Singer. Cognitive, social, and psychological determinants of emotional state. *Psychological Review*, 69:379–399, 1962.

Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656, July and October 1948.

Herbert A. Simon. Motivational and emotional controls of cognition. *Psychological Review*, 74:29–39, 1967.

Aaron Sloman. Varieties of affect and the CogAff architecture schema. In Colin Johnson, editor, *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing*, pages 39–48, Heslington, York, England, March 2001. AISB.

Richard S. Sutton and Andrew G. Barto. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88(2):135–170, 1981.

Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211: 453–458, 1981.

R. M. Yerkes and J. D. Dodson. The relation of strength of stimulus to rapidity of habit formation. *Journal of Comparative Neurology and Psychology*, 18:459–482, 1908.

R. Zajonc. Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35:151–175, 1980.