

Using Anticipation to Create Believable Behaviour *

Carlos Martinho

Instituto Superior Técnico
Av. Prof. Cavaco Silva, TagusPark
2780-990 Porto Salvo, Portugal
e-mail: carlos.martinho@tagus.ist.utl.pt

Ana Paiva

INESC-ID
Rua Alves Redol 9, Apartado 13069,
1000-029 Lisboa, Portugal
e-mail: ana.paiva@inesc-id.pt

Abstract

Although anticipation is an important part of creating believable behaviour, it has had but a secondary role in the field of life-like characters. In this paper, we show how a simple anticipatory mechanism can be used to control the behaviour of a synthetic character implemented as a software agent, without disrupting the user's suspension of disbelief. We describe the emotivevector, an anticipatory mechanism coupled with a sensor, that: (1) uses the history of the sensor to anticipate the next sensor state; (2) interprets the mismatch between the prediction and the sensed value, by computing its attention grabbing potential and associating a basic qualitative sensation with the signal; (3) sends its interpretation along with the signal. When a signal from the sensor reaches the processing module of the agent, it carries recommendations such as: "you should seriously take this signal into consideration, as it is much better than we had expected" or "just forget about this one, it is as bad as we predicted". We delineate several strategies to manage several emotivevectors at once and show how one of these strategies (meta-anticipation) transparently introduces the concept of uncertainty. Finally, we describe an experiment in which an emotivevector-controlled synthetic character interacts with the user in the context of a word-puzzle game and present the evaluation supporting the adequacy of our approach.

Introduction

Artificial Intelligence has long sought to construct autonomous creatures. The thought of these entities brings special delight when they are imagined to project a sense of being "really there". This sensation gave birth to the concept of *believability*. Although the notion falls prey to its subjective nature, more than two decades of research have followed the seminal definition forwarded by Bates (1994), outlining a believable character as one able to maintain the user's "suspension of disbelief". Even the most unrealistic behaviour of a synthetic character must be consistent with the expectations created in the user through the character's presentation. As such, *anticipation* plays an important part in the creation of believability.

*This work is supported by EU project MindRACES: from reactive to anticipatory cognitive embodied systems (IST-511931). Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

The quest for believability has sent researchers on two different paths: a pragmatical approach inspired in Arts such as drama and character animation (Thomas & Johnston 1994), and another that strives for higher levels of autonomy by providing the synthetic creation with biologically plausible (Blumberg 2002) or psychologically sound behaviour (Gratch & Marsella 2003). Both paths emphasize the concept of believability as a dimension of synthetic performance closely related to the adequate expression of *emotion*. As a result, most believable characters have some affective model (Picard 1997) underlying their behaviour.

Several models of emotions have been proposed to aid achieving believability. However, none explicitly integrates the concept of *anticipation* in the creation of life-like behaviour. Anticipation is usually found diluted in the planning mechanisms of the synthetic character (Gratch & Marsella 2003) or disguised as an emotion by itself, one that involves pleasure in considering some expected or longed-for good event, or irritation at having to wait, as in Plutchik's theory of emotions (Plutchik 1991). Anticipation has had but a secondary role in the creation of believable characters.

In this paper, we show how a simple anticipatory mechanism — that we call *emotivevector* — can lead to believable behaviour. We describe the emotivevector, how it monitorizes the sensor history of the agent, anticipates the next sensor state and extracts salience and qualitative information from the mismatch between prediction and the sensed value. Unlike approaches such as Breazeal and Scassellati's motivation driven perceptual maps (Breazeal & Scassellati 1999), we compute salience by anticipating the sensor state rather than by just reacting to it.

We delineate strategies to manage several emotivevectors at once and show how meta-anticipation can represent the concept of uncertainty. Finally, we describe an experiment where the information produced by the emotivevector was used to control a synthetic character and show how the evaluation asserted the adequacy of the approach to produce believable behaviour.

Anticipation and Emotion

Rosen (1985) defined an anticipatory system as one that contains a predictive model of itself and/or its environment, which allows it to change state at an instant in accordance with the model's predictions pertaining to a latter instant.

Anticipatory systems possess two interesting features (Martinho *et al.* 2005). First, *surprise* is an inherent characteristic of anticipatory systems, as purely reactive systems cannot be surprised. Surprise essentially arises as a mismatch between what is expected to be perceived and what is actually perceived. Second, the basal functionality of an anticipatory system relies on its ability to predict whether the system is going in the “desired” or “undesired” direction, which is but the simplest form of valenced affective reaction.

The idea behind the concept of emotivector is to capture these two features in their most basic form, and assess their potential in the creation of believable behaviour.

Emotivector

The emotivector is an anticipatory mechanism attached to a sensor. Each emotivector is associated with a one-dimensional aspect of perception, and all emotivectors are kept together in a salience module, responsible for their management as a whole. Figure 1 shows the agent architecture.

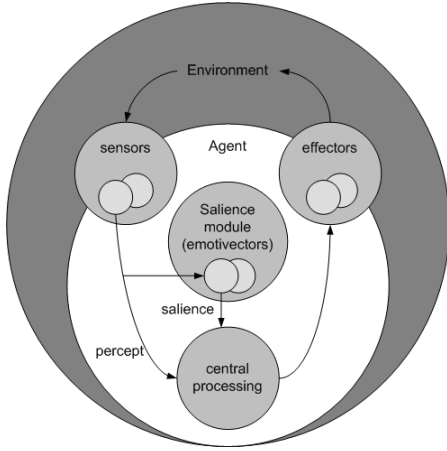


Figure 1: Agent Architecture

When a percept is sent to the processing module, the associated emotivector catches its value and performs the following steps: (1) using the history of percepts, the emotivector computes the next expected value of the sensor, relying on a hybrid algorithm based on the Kalman filter and the generalized recirculation algorithm; (2) by confronting the expectation with the actual sensed value, and using a model inspired in the psychology of attention, the emotivector computes a preliminary salience for the percept, and; (3) a sensation is generated according to a model inspired in the psychology of emotion. The combination of both “attentional” and “emotional” interpretation is added to the percept. When a percept reaches the processing module, it carries recommendations such as “you should seriously take this signal into consideration, as it is much better than we had expected” or “just forget about this one, it is as bad as we predicted”

The following sections describe each step. For the sake of simplicity, all sensor values are normalized to $[0, 1]$.

Percept Salience (Attention Model)

The computation of an a-priori salience for a percept is based on a model of attention inspired by Posner’s exogenous (automatic reflexive control) and endogenous (voluntary control) systems (Posner 1980) and how these systems interact according to Müller’s hypothesis (1989).

Our model of attention is implemented as follows. At time $t-1$, the emotivector value is $x_{t-1} \in [0, 1]$. Using its history at time $t-1$, the emotivector estimates a value for next time t (\hat{x}_t), and predicts that its value will change by $\Delta\hat{x}_t = \hat{x}_t - x_{t-1}$. At time t , a new value is sensed (x_t), and a variation $\Delta x_t = x_t - x_{t-1}$ is actually verified (see Figure 2).




Figure 2: Computing Salience

The exogenous component at time t (EXO_t) is based on the estimation error and reflects the principle that the least expected (i.e. more surprising in terms of mismatch) is more likely to attract the attention:

$$EXO_t = (x_t - \hat{x}_t)^2$$

If no further information is given to the emotivector, the exogenous component is the only factor contributing to salience. However, if a “searched” value (s_t) exists at time t , the endogenous component (END_t) of the percept is computed as:

$$\begin{aligned} \Delta s_t &= (x_t - s_t)^2 \\ \Delta \hat{s}_t &= (\hat{x}_t - s_t)^2 \\ END_t &= \Delta \hat{s}_t - \Delta s_t \end{aligned}$$

END_t reflects whether the change in the distance to the “searched” value is better or worse than the expectation and the level of mismatch (see Figure 2). Unlike EXO_t , which is always nonnegative, END_t is valenced: an increase in the expected search distance is assumed negative, while a reduction is modeled by a positive value.

The combination of both exogenous and endogenous components define the salience of the percept. However, an emotivector with a search value can also provide a qualitative interpretation of the percept, as explained in the next section.

Percept Sensation (Emotion Model)

Our affective model is inspired by early behavioural theories and models of emotion.

We assume that emotions are conditioned responses of primary sensations (Harlow & Stagner 1933) and concentrate our model in the generation of these sensations. Emotions per se are left to be handled by the processing module of the agent.

Sensations are defined across two dimensions (Young 1961): valence and change. As such, our model considers four basic sensations: positive increase, positive reduction, negative increase and negative reduction. The emotivevector estimation is used to anticipate a reward or punishment which, when confronted with the actual reward or punishment (Hammond 1970), triggers one of the four basic sensations. The intensity of each emotion is given by the endogenous component of the percept.

Thus, our model considers the following four primary sensations¹:

S+ (positive increase) If reward is anticipated and the effective reward is stronger than the expected, an S+ sensation is thrown.

\$+ (positive reduction) If reward is anticipated but the effective reward is weaker than the expected, a \$+ sensation is thrown.

S- (negative increase) If punishment is anticipated and the effective punishment is stronger than expected, an S- sensation is thrown.

\$- (negative reduction) If punishment is anticipated but the effective punishment is weaker than expected, a \$- sensation is thrown.

Following our nomenclature, the expected reward (\hat{R}_t) for time t and the sensed reward (R_t) at time t are computed as follows:

$$\begin{aligned}\hat{R}_t &= \Delta s_{t-1} - \Delta \hat{s}_t \\ R_t &= \Delta s_{t-1} - \Delta s_t\end{aligned}$$

Table 1 shows the eliciting conditions for each basic sensation. Each sensation has an intensity defined by the endogenous component END_t .

predicted	sensed	sensation
$\hat{R}_t > 0$	$R_t > \hat{R}_t$	S+
$\hat{R}_t > 0$	$R_t < \hat{R}_t$	\$+
$\hat{R}_t < 0$	$R_t < \hat{R}_t$	S-
$\hat{R}_t < 0$	$R_t > \hat{R}_t$	\$-

Table 1: Computing Sensation

¹We follow Millenson's symbolism to name our sensations, as symbols are not connoted with an exact word which, by itself, would imply a certain intensity.

Percept Prediction (Anticipatory Model)

The computation of the emotivevector salience relies on the capacity of the emotivevector to predict its next state. Before anything else, we define the model that we expect the sensed data to follow – if the signals are totally random, no prediction strategy can be evaluated for adequacy.

Model

As there is no a-priori knowledge of the signal, we follow a simple assumption: that the intensity i of a signal will change by a random small amount $\Delta i \in [-\epsilon, \epsilon]$ at each discrete time step (defined by the sensor rate), for a random time slice $\Delta t \in]0, \Delta t_{max}]$, before suddenly changing to a random new value in the interval $[0, 1]$. In other words, the sensed value will tend to remain constant except for certain points in time.

Kalman Filtering

A possible estimator for such a signal is the Kalman filter (Kalman 1960), a set of mathematical equations that provides efficient computational recursive means to estimate the state of a process in a way that minimizes the mean of the squared error.

A Kalman filter estimating a random constant was implemented. The filter is composed of two sets of equations: the *time update* equations, that project forward (in time) the current state and error covariance estimates to obtain the a-priori estimate for the next time step and; the *measurement update* equations, that are responsible for the feedback i.e. for incorporating a new measurement into the a-priori estimate to obtain an improved a posteriori estimate. In the following equations, R is the measurement noise covariance and Q is the process noise covariance:

1. Time update equations

(a) Project the state ahead: $\hat{x}_t^- = \hat{x}_{t-1}$

(b) Project the error covariance ahead: $P_t^- = P_{t-1} + Q$

2. Measurement update equations

(a) Compute the Kalman gain: $K_t = P_t^- / (P_t^- + R)$

(b) Update estimate: $\hat{x}_t = \hat{x}_t^- + K_t(x_t + \hat{x}_t^-)$ (1)

(c) Update the error covariance: $P_t = (1 - K_t)P_t^-$

The need for parameters Q and R to be initially fine tuned suggested another path.

Simple Predictor

We rewrote \hat{x}_t from Equation 1 as:

$$\hat{x}_t = x_{t-1}^- \frac{R}{P_{t-1} + Q + R} + x_t \frac{P_{t-1} + Q}{P_{t-1} + Q + R}$$

and then made the following substitutions:

$$\alpha = \frac{P_{t-1} + Q}{P_{t-1} + Q + R}, \beta = \frac{R}{P_{t-1} + Q + R}$$

Noticing that $\alpha + \beta = 1$, we rewrote \hat{x}_t as:

$$\hat{x}_t = \hat{x}_{t-1}(1 - \alpha) + x_t \alpha$$




Figure 3: Simple Predictor

Figure 3 shows this simple predictor graphically. The predictor is a function of the estimated value and the sensed value, both competing in the contribution to the prediction.

This result resembles to the generalized recirculation algorithm (Hinton & McClelland 1988), an algorithm that enabled back-propagation to be implemented in a more biologically plausible manner, by adopting two activation phases: the *minus phase*, where the outputs of the system represent the *expectation* of the system, as a function of the standard activation settling process in response to a given input pattern; the *plus phase*, where the environment is responsible for providing the *target* output activation. In this algorithm, the learning is essentially the delta rule (Widrow & Hoff 1960) confronting the expectation with the sensed value.

Similarly, we used the delta rule with a learning rate set to the endogenous salience, in such a way that stimuli sensed under strong sensations are more relevant in terms of prediction. Hence, $\Delta\alpha_t$ is computed as follows:

$$\Delta\alpha_t = \xi_t(x_k - \hat{x}_{k-1})\alpha_t$$

where $\xi_t = \text{END}_t^2$.

Even if this predictor is not as optimal as the ones it is inspired by, it provides a good efficiency/adaptation relation that performs well in real-time over unpredictable signal dimensions, and does not require any previous fine tuning to work.

Emotivector Management

In this section, we discuss how the salience module selects among all percepts which are *a-priori* more relevant to the processing module of the agent.

Winner-Takes-All

The most immediate approach is to use a *winner-takes-all* strategy, and select only the percept with the highest salience.

Although this may be a convenient approach when processing power is a crucial factor, the highest percept can easily *hide* other percepts, which are just “a little less salient”, but in fact more important to the processing module.

Salience Ordering

To allow for more than one percept to reach the processing module at the time, an alternative strategy is *salience ordering*: all percepts are ordered by salience and sent to the

processing module. The central processing module will then attend to the most relevant percept, then the second, and so on, until no more processing resources are available.

Although using the full processing power of the system on which it is built, the agent will usually spend most of its processing time dealing with percepts of very low significance and waste precious processing power unnecessarily, which is something that one wants to avoid as to allow for other more demanding processes (e.g. graphics and physics) to perform at their best.

Of course, a minimum threshold could be set, below which the percepts would not be sent for processing. The other percepts would continue to be processed by salience order. However, how is the threshold value chosen? It depends on the setting context and may vary dynamically according to the situation. Another approach would be to only process the first N most relevant percepts, but again, we would fall into the same fallacy.

In the end, this strategy may work, but fine-tuning would make it inelegant.

Meta Anticipation

Our approach is to use a meta-predictor, whose role is to assess the salience of its associated predictor. The meta-predictor is perfectly identical to all other predictors. The only difference is that the meta-predictor receives prediction errors, while the other predictors receive percepts.

Each time a new value is sensed, there is usually a mismatch between the prediction and the sensed value. This prediction error is fed to the meta-predictor. Based on the history of prediction errors, the meta-predictor computes the next expected error, using the same mechanism the average predictor uses to compute its next prediction. When a new value is sensed, the prediction error is computed and sent to the meta-predictor that will compare it with its error expectation. If the prediction error is higher than the meta-predictor estimated error, the percept sensed by the emotivector is marked as relevant to the processing module. Otherwise, the percept is marked as non-relevant.

This emotivector management strategy has the advantage of not requiring parameter fine-tuning, and was chosen for our experiment.

Uncertainty

An interesting side-effect of meta-prediction is that it provides an error margin for estimation. In other words, when a predictor estimates a value, the meta-predictor estimates how trustworthy the prediction is. As such, meta-prediction models the concept of *uncertainty*.

Uncertainty allows us to add an “as expected” parameter to our model of sensations, transforming our initial set of four sensations (S+, \$+, S- and \$-) into a full set of nine basic sensations, where sensations as “this percept is as bad as we expected it to be” allow the enrichment of the vocabulary of sensations associated with a percept. Table 2 shows the relation between the nine-sensation model and the four-sensation model presented earlier.

expected	sensed		
	more R	as expected	more P
reward (R)	stronger R (S+)	expected R	weaker R (S+)
negligible	unexpected R	negligible	unexpected P
punishment (P)	weaker P (S-)	expected P	stronger P (S-)

Table 2: Extended Sensation Model

Experiment

Believability is a difficult concept to experimentally assess. It is usually measured through questionnaires evaluating the satisfaction of the subjects, and relating the satisfaction level to the suspension of disbelief provoked by the experiment.

Although a questionnaire is invaluable, we designed an interactive task where the user success depends on the consistent and believable behaviour of a synthetic character, in such a way that the observation of the task itself provides less subjective measurements regarding believability.

We created a word puzzle game in which Aini, a synthetic flower built with our emotivevector model, helped the user to uncover a four-letter word solely by reacting to her actions. To construct the word, the player must use a set of letter-boxes, each one with an alphabet letter drawn onto it. The word is constructed by placing the boxes onto wooden platforms representing the word letters and their relative position in the word, a tridimensional version of the hangman dashes. Only some of the letter-cubes are relevant, others act as distractors. To allow for a more intuitive interaction with the letter-boxes, the simulation is fully tridimensional and physics-based. Figure 4 shows the applicational setting.



Figure 4: Experiment Application

Emotivevector Control

Our first use of the emotivevectors was to control Aini’s focus of attention, a task that only required the exogenous component. We provided Aini with a 5×5 grid of sensors (and associated emotivevectors) measuring the distance to the nearest objects in Aini’s field of view. The strongest salience designates the target of Aini’s gaze and the intensity controls the speed of movement. This approach automatically implemented the “casual look around” as well as the “quickly look at something new” behaviours.

Our second use of the emotivevector was to monitor the progress of the word completion task. A single emotivevector was regularly fed with a measure of completeness of the task. A search value, corresponding to the task being completed, activated the endogenous component. When salient, the sensations created by the emotivevector triggered different expressive behaviours.

Consider the example of a user trying to uncover the word “LIFE”. Initially, Aini is in a neutral state. Because the user has now been playing for quite a while without trying any letter, the margin of prediction error has dropped, and Aini is pretty sure nothing is going to happen.

Suddenly, the user places her first guess (“L”) in a wrong placeholder. As a result, Aini expresses an “unexpected P” sensation. It stops moving and lowers its head. Aini expects more punishment to come but is not sure about it. Seeing the negative reaction of the flower, the user removes the letter “L” from the word. As the completeness increases by a significant amount, Aini expresses a “weaker P” sensation: it rises its head and waves it encouragingly. Notice that if the user placed the “L” in the same wrong place again, Aini would react differently from the first time, as the conditions would have changed.

Now, the user places the “L” in the correct position. As it is still under great uncertainty, an “expected R” sensation is triggered: the agent now looks confident about the user capabilities. Afterward, the user places the “I” in the right position. As the margin of error now diminished, the value is outside the predictive margin and launches a “stronger R” sensation, expressing total bewilderment.

This example shows the richness of situations that can be obtained from the use of the emotivevector basic sensation alone, and how perception history and timing triggers different basic sensations in response to the same user’s action.

Evaluation

To evaluate our model, we asked users to play the word puzzle with four different synthetic characters (presented in different orders) sharing the same appearance but having distinct behaviours. Two control characters, with idle and random behaviour respectively, confirmed that the task is impossible to complete without consistent help. Two other characters evaluated the adequacy of our model: one used our emotivevector control; the other implemented an approach used by the current generation of computer games, such as *Fable*² and *Oblivion*³. In our game, it translates to: look at the nearest object or latest word change and use a positive or negative expression depending on action appraisal.

After interacting with each character, each subject answered a questionnaire evaluating the following aspects: the smoothness of the animation, the naturalness of the behaviour, the level of expectation of the subject regarding the provided behaviour, the similarity with the behaviour of a person in the same situation, the ease of understanding the displayed attention and emotion, and the general impression left by the synthetic character.

²Fable: The Lost Chapters, Lionhead Studios, 2005.

³The Elders Scrolls IV: Oblivion, Bethesda Softworks, 2006.

A total of 280 puzzles were played by both male and female subjects from 5 to 79 years-old, with different levels of familiarity with computers. No subject finished the game with the control characters: although it would be possible to use a brute force approach, the presence of an inconsistent character led the user to quit. All subjects managed to finish the task with the emotivevector-based character, but only 20.6% of the subjects were successful with the game-based approach. Unexpectedly, the game-based behaviour induced the feeling of “being cheated on” and made the user quit. Only users experimenting with one letter at a time, while all other letters were kept far away, were successful.

To reach a better understanding of the results, we performed a 2-dimension HOMALS on the data from the questionnaires, that accounted for 73.6% of the variance. From the HOMALS quantification, two variables were identified as totally unrelated to believability: “smoothness” and “previous knowledge of the word”. A Fisher test strengthened (exact sig 0.678) that the real challenge is to understand the expression of the synthetic character. Indeed, even the younger subjects were able to discover unknown words. 93.7% of subjects rated “guessing a known word without help” as very difficult (5 in 5 degree Likert scale) while only 6.4% of the subjects think that “finding an unknown word with help” is difficult (4 or 5 in 5 degree Likert scale).

The HOMALS first dimension stated that a subject who succeeded at the task generally enjoyed the experience and found that attention and emotion were easy to perceive. The second dimension suggests that when a subject evaluates emotion recognition or task difficulty as “extreme” (in the sense of very good or very bad) this evaluation is reflected in a more “extreme” assessment of believability.

During the experiments, two unexpected results were found. First, the exogenous component based synthetic “vision” provided the user with a natural way to interact with the agent (e.g. by waving a box to attract attention to it). Second, the endogenous component of the task-completeness emotivevector allowed a rich and non-repetitive behaviour of the synthetic character that seemed to account for the past history of interaction in a meaningful manner.

Believability seems to be influenced by all evaluated aspects of the experience except for the smoothness of the animation. It is congruent with the fact that more complex animation does not necessarily means more believable. Our results confirm that both the focus of attention and emotional state play a more important role in the definition of believability.

Conclusion and Future Work

In this paper we discussed how a simple anticipatory mechanism, the emotivevector, can be used to produce believable behaviour. We described the emotivevector mechanism, inspired by theories of attention and emotion, and how it was successfully used in a concrete interactive experiment.

Our model does not intend to substitute any current approach but rather complement the agent architecture for life-like characters: the salience module exists “outside” the basic agent components, and as such, can be added or removed

from an agent with a low cost. Another approach for emotivevector integration in more complex models is to substitute each sensor by a “sensor+emotivevector” combination. We are currently using this approach to integrate emotivevectors in agents with higher-level cognitive theories of emotion and anticipation implemented at the processing level to model *cautiousness*.

References

- Bates, J. 1994. The role of emotion in believable agents. Technical report, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- Blumberg, B. 2002. *Exploring Artificial Intelligence in the new Millennium*. Morgan Kaufmann Publishers. chapter D-Learning: What learning in dogs tells us about building characters that learn what they ought to learn.
- Breazeal, C., and Scassellati, B. 1999. A context-dependent attention system for a social robot. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*.
- Gratch, J., and Marsella, S. 2003. Modeling coping behavior in virtual humans: Don’t worry, be happy. In *Proceedings of Second International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press.
- Hammond, L. 1970. Conditioned emotional state. In *Physiological Correlates of Emotion*. Academic Press.
- Harlow, H., and Stagner, R. 1933. Theory of emotions. *Psychological Reviews*.
- Hinton, G., and McClelland, J. 1988. Learning representations by recirculation. In Anderson, D., ed., *Neural Information Processing Systems*. American Institute of Physics.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering* 82(Series D):35–45.
- Martinho, C.; Micelli, M.; Dias, J.; Paiva, A.; and Castelfranchi, C. 2005. Anticipation and emotion. Technical report, MindRACES Technical Report.
- Muller, H., and Rabbit, P. 1989. Reflexive orienting of visual attention: time course of activation and resistance to interruption. *Journal of Experimental Psychology: Human Perception and Performance*, (15).
- Picard, R. 1997. *Affective Computing*. MIT Press.
- Plutchik, R. 1991. Emotion and evolution. In Strongman, K. T., ed., *International review of studies of emotions, vol 1*. Chichester: Wiley.
- Posner, M. 1980. Orienting of attention. *Quarterly Journal of Experimental Psychology* (32).
- Rosen, R. 1985. *Anticipatory Systems*. Pergamon Press.
- Thomas, F., and Johnston, O. 1994. *The Illusion of Life*. Hyperion Press.
- Widrow, B., and Hoff, M. 1960. Adaptive switching circuits. *IRE WESCON Convention Record*.
- Young, P. 1961. *Motivation and Emotion*. Wiley.