Decaying Confidence Functions for Aging Knowledge in Expert Systems

Joseph Pasquale

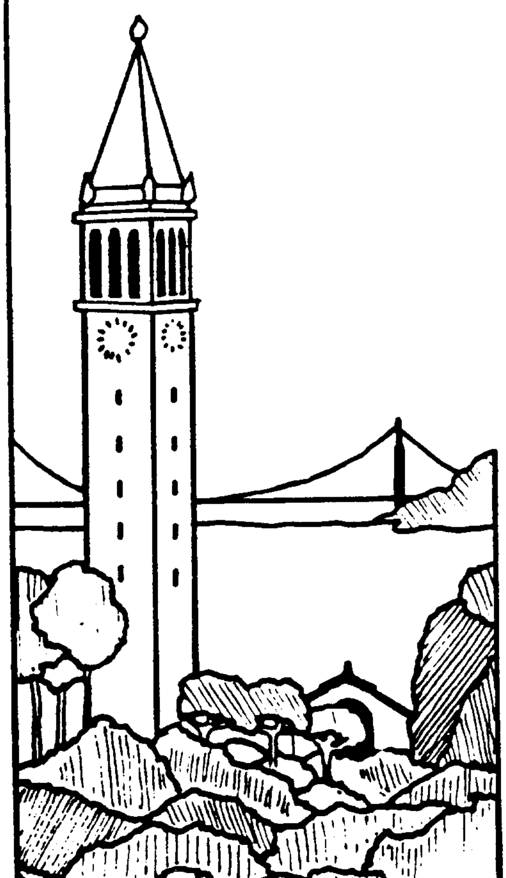

Report No. UCB/CSD 87/335

January 1987
PROGRES Report No. 86.9

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Decaying Confidence Functions for Aging Knowledge in Expert Systems

Joseph Pasquale

Computer Systems Research Group
Computer Science Division
Department of Electrical Engineering and Computer Sciences
University of California
Berkeley, CA 94720

Abstract

Expert systems which maintain knowledge about objects whose attributes are time-variant must have an awareness of time. This awareness can be made manifest by incorporating time in the quantification of uncertainty of aging knowledge about such objects.

Many expert systems use some method to quantify the degree of belief, or uncertainty, of their knowledge. Examples of these methods include Bayesian probability theory, certainty factors of EMYCIN, the Dempster-Shafer theory, and fuzzy set theory. These methods offer different representations for measures of confidence, and different calculi for combining these measures. We describe an extension to such confidence measures by adding a dimension of time.

We propose the concept of Decaying Confidence Functions to express the time-varying uncertainty of aging knowledge. Decaying confidence functions specify how confidence in knowledge should decrease as the knowledge gets older. We describe how this can lead to efficiencies in expert systems which must deal with time-varying information, such as expert systems used for monitoring real-time systems.

This work was supported by an I.B.M. Fellowship and by the Defense Advanced Research Projects Agency (DoD), monitored by the Naval Electronics Systems Command under contract No. N00039-84-C-0089. The views and conclusions contained in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied, of I.B.M., the Defense Research Projects Agency or of the US Government.

1. Introduction

Items in the knowledge base of an expert system (i.e., facts, hypotheses, propositions, and sometimes even rules) are generally believed with varying degrees of uncertainty. To express this uncertainty, many expert systems associate a measure of confidence with each item of knowledge. The value of qualifying information with confidence measures has been recognized ever since the pioneering work of MYCIN [1], which used "Certainty Factors" to express uncertainty of propositions. Today, there are many methods in use: Bayesian probability methods [2], certainty factors [3], Dempster-Shafer Theory of Belief Functions [4], fuzzy logic [5], and other multi-valued logics [6]. These methods differ in how uncertainty is represented, such as the point probabilities used by Bayesian methods, the intervals of uncertainty used by Dempster-Shafer theory, and the set membership quantifications used by fuzzy logic. We refer to these different representations generically as *confidence measures*. These methods also differ in how their respective confidence measures are combined when chains of inferences take place.

When a confidence measure is assigned to some item of knowledge, it generally does not change unless there exists evidence to cause it to change. Yet, there are times when it makes sense to change a confidence measure, even though no new evidence *explicitly* appears. This occurs in systems where information about objects is *time-variant*. An intelligent agent, such as an expert system, might detect that a given object was in a certain state at a particular time. It would be highly confident of this knowledge at that time, if, for instance, it was able to detect the object's state directly using sensors. As time passes though, the agent expects the object's state to change and therefore lowers its confidence that its original knowledge is still correct. One might say that *implicit* evidence, the knowledge of the time-variance of object states, caused the confidence measure to change.

This scenario is very common to expert systems used for monitoring *dynamic systems*, environments comprised of objects whose attributes are time-variant. In fact, these environments typically have many sources of time-varying information which must be monitored and correlated in real-time, making the problem very difficult. As excellent examples, Steeb et al. [7] describe two problems, Air-traffic Control and Remotely Piloted Vehicle Fleet Control, where expert system technology is being applied to monitor and control highly dynamic real-time systems.

We argue that expert systems used for monitoring dynamic systems must be aware that some of their knowledge is based on time-sensitive information. Building into expert systems knowledge about the process of information becoming obsolete, and that this process can be gradual, abrupt, or some degree in between, can lead to more efficient ways of monitoring and interpreting what is happening in the observed system. Consequently, our approach will be practical (as opposed to a more theoretical approach such as in *temporal logic*) and will focus on providing an efficient mechanism for dealing with time-varying uncertainty.

In this paper, we introduce the idea of confidence measures which decay over time, which we call **Decaying Confidence Functions (DCFs)**. The concept of DCFs is not a new way of expressing uncertainty, nor is it a way of combining confidence measures. Rather, it is a generalization of all currently available confidence measures through the addition of a dimension of time. The use of the DCF concept is essentially governed by the nature of the an expert system's knowledge, specifically that this knowledge is based on information which varies over time. In section 2, we describe formally the concept of DCFs. In section 3, we describe conditions which warrant the use of DCFs. In section 4, we describe what gains in expert system efficiency are derived from using DCFs. Finally, we summarize our conclusions in section 5.

2. Definition of Decaying Confidence Functions

In many expert systems, a confidence measure is assigned to a basic unit of knowledge, such as a fact or hypothesis. For example, in EMYCIN, these units are called attribute-object-value triples or contexts [8], and the confidence measure is called a certainty factor. We are most interested in knowledge based on time-varying information, so let us define a basic unit of knowledge parameterized by time in the following way.

Let

$$K_X(t) = (A_1 = a_1, A_2 = a_2, \cdots) \tag{1}$$

be a unit of knowledge comprising information about the values of object X's attributes $(A_1, A_2, \cdots)$ at time t. We will say that this knowledge was established at time t, or that the knowledge is about information acquired or detected at time t.

For example, consider a ship S with attributes (Location, Speed, Direction). A realization of $K_X(t)$ for ship S at some specific time t_s would be:

$$K_S(t_s) = (\begin{array}{lcl} \text{Location} & = & < 7^\circ 54'34'' \text{ N}, 31^\circ 20'16'' \text{ W}>, \\ \text{Speed} & = & < 60 \text{ nautical m.p.h.}>, \\ \text{Direction} & = & < \text{directly East} >) \end{array}$$

What can we say about the utility of this knowledge as a function of time. After one second, the ship would travel a distance of approximately 1" making part of the information given by $K_S(t_s)$, namely the location, invalid for the current time. To be precise, by "invalid for the current time" we mean that the value of the location of the ship as given by $K_S(t_s)$ will not be the same value for the location of the ship at the time t_s + 1 second. In fact, we would expect all the information given by $K_S(t_s)$ to eventually become invalid some time after t_s. Does this mean $K_S(t_s)$ is virtually useless to the expert system? Must the expert system continually receive updates (at least one per second) if it wants to make use of information about the ship?

Of course, the answer to these questions is no. Information, which describes something at one point in time, which then does not apply at some future time, is not useless information in future times. On the contrary, it is very useful for making inferences about the future.

Inference in monitoring systems usually comes under the guise of estimation, forecasting or prediction. Typically, the environment to be monitored is sampled periodically. Information about the environment at some future time since that last sample is *predicted*, and is usually based on previous samples. These predictions, just like information received directly from the environment, are not believed with complete certainty. We will treat these predictions simply as knowledge derived from other knowledge which was based on directly sensed information (i.e., knowledge as defined by (1)), and represent it as follows.

Let

$$K_X(t \mid \{K{:}t_r\}) = (A_1 = a_1, A_2 = a_2, \cdots) \tag{2}$$

be knowledge comprising information about the attributes of object X, given some set of previously established knowledge, specified by $\{K{:}t_r\}$ which is a condensed notation for

$$\{K(t_1), K(t_2), \cdots : \max(t_1, t_2, \ldots) = t_r \leq t\}$$

Note that the most recent knowledge, used as a basis of the derived knowledge, was established at the time t_r which is defined less than or equal to time t.

The procedure used for inferring the future, or predicting, may be elaborate or very simple, and its ability to predict correctly may be correspondingly good or bad. As a measure of goodness, let

$$D(\ Kx(t \mid \{K{:}t-\tau\}),\ Tx(t)\ ) = r,\ r \geq 0,\ \tau > 0 \tag{3}$$

be a measure of the difference between $Kx(t \mid \{K{:}t-\tau\})$ and $Tx(t)$, the true values of the attributes of object X at time t. τ is the age of the most recently established knowledge comprising information used as a basis for Kx. The measure is quantified by the non-negative real number r.

Finally, we make the following reasonable assumption that, in general, as the time between the prediction and the age of the knowledge on which the prediction is based becomes larger, we expect the prediction to become less accurate.

We can express this as an inequality of expected values:

$$E[\ d(t+\xi,\ t-\tau)\ ] > E[\ d(t,t-\tau)\ ],\ \tau>0,\ \xi>0 \tag{4}$$

where

$$d(t,\ t-\tau) = D(\ Kx(t \mid \{K{:}t-\tau\}),\ Tx(t)\ )$$

and again, τ is the age of the most recent knowledge used as a basis for Kx.

Inequality (4) simply states that we expect the difference between the predicted values and the true values of attributes of an object to get larger as the interval between the time corresponding to the prediction and the time of most recent knowledge used to base the prediction gets larger. This inequality implies that the more recent our knowledge about something is, the more useful it is. Therefore it is important to be aware of its age. The goal is to quantify this usefulness of knowledge as a function of time.

Let us now define a **Family of Decaying Confidence Functions**:

$$C_\Theta(K,\tau) = \langle \mathit{confidence\ measure} \rangle \tag{5}$$

K is a time-varying unit of derived knowledge as defined in (2)

τ is a time interval specifying the age of the knowledge used as a basis for K

Θ is a set of parameters which select one of a family of functions defined by C

The range of $C_\Theta(K,\tau)$ is a confidence measure, such as a certainty factor, point probability, probability range. The concept of Decaying Confidence Functions is applicable to all confidence measures; it focuses on the role of age of the information being qualified by the confidence measure.

Equation (5) defines a family of functions of a common form, and Θ defines the space of these functions. The parameters comprised by Θ define function characteristics, such as initial value, rate of decay, and points of discontinuity. To make clear the role of Θ, let us consider the following example.

For simplicity, let us assume our confidence measure of choice is the point probability. For a physical system, such as ship S described in the previous example, we might find the following general decaying confidence function appropriate:

$$C_\Theta(Ks,\tau) = \rho^{-\upsilon\tau} \tag{6}$$

$\Theta = (\rho,\ \upsilon)$
ρ = prediction factor ($\rho \geq 1$)
υ = measure of variability ($\upsilon \geq 0$)

Equation (6) defines a family of decaying exponential functions. The range of these functions is a real number between 0 and 1, a point probability. The space of this family of functions is defined by two parameters, ρ and υ. Parameter ρ is a measure of the expert system's *predictive* capabilities. Given $Ks(t_s)$, for some fixed time t_s, how well can we predict $Ks(t \mid \{Ks{:}t_s\})$ where the age is given by $\tau = t - t_s$. The better our prediction algorithm, the smaller ρ should be, causing (6), the measure of confidence that $Ks(t \mid \{Ks{:}t_s\})$ is true for future times, to decay at an overall slower rate.

To describe parameter υ, let us define a history of $Ks(t)$ to be a set of time-varying knowledge units $Ks(t)$ at particular times $t_1, t_2, \ldots, t_i < t$ (i.e., $\{Ks(t_1), Ks(t_2), \cdots\}$). Parameter υ is a measure of the *variability* of the history of $Ks(t)$. This measure might be weighted toward more recent values of $Ks(t)$ in the history. The expert system could compute υ itself from whatever history it has of $Ks(t)$. Or, it may learn the value of υ by the ship transmitting current values of υ every so often. Consequently, the more variable the past history of $Ks(t)$ is, the faster our confidence decays of $Ks(t \mid \{Ks{:}t_s\})$ remaining true as t increases.

Notice that υ is a parameter which may be updated periodically in real-time, consequently updating the confidence function for a particular item of information. In contrast, parameter ρ is a relatively static value, since the expert system's prediction algorithm would not be expected to change as often. Hence the reason why there are two distinct parameters, even though they have the same effect on the confidence function.

This example illustrates how a particular family of decaying confidence functions defines a sequence of point probabilities which decay in time, based on the age of $Ks(t_s)$, the expert system's predictive powers, and on the variability of a past history of $Ks(t)$.

3. Conditions for using Decaying Confidence Functions

Decaying Confidence Functions are a generalization of current confidence measures in that they add a dimension of time to current methods. The use of DCFs is justified in certain types of environments; in others, it is simply an added complexity which is unnecessary. In this section, we describe an environment where we have found the use of DCFs to be beneficial. More importantly, we generalize our observations from our experiences to describe conditions warranting the use of DCFs.

We are using DCFs in an expert system which is used to monitor and manage a distributed computer system [9]. The expert system is itself fully distributed, and one of its functions is to determine the states of shared resources comprising the distributed computer system. It then makes decisions concerning scheduling and allocation of these resources. A major problem is that of obtaining an up-to-date, correct view of the global state of all resources, which must essentially be inferred using past information. It is therefore important that the expert system qualify all knowledge about the states of remote resources with the time that knowledge was established.

Thus, our first condition for using DCFs is that the expert system must be dealing with information which is time-sensitive. Specifically, the value of knowing the information decays as time elapses from the point of learning the information. This is in contrast, for example, to dealing with information which is assumed to always be true, thus insensitive to time.

In our experimental distributed system as in most systems, communication is not free. Communication costs are contributed in varying degrees by the source of information, the transmission channels, and the receiver. For example, if a source must do considerable preprocessing of information before it is sent, or if, because of protocol limitations, it must send in a serial fashion to many destinations, the sender will incur a significant cost. There is significant cost associated with simply using a transmission channel because of the incurred delays, especially when information must be distributed to all nodes of a large complex network. Finally, receiving information would be considered expensive in cases where considerable processing of the received information must take place, such as the processing of images. The point is that since communication incurs a non-trivial cost, it is worth minimizing if possible.

We generalize this observation by proposing this second condition, which is that DCFs are useful where it is expensive to communicate information which must somehow be learned. Given that communication is costly, it might actually be worth using old information to infer current reality if possible rather than to acquire updated information. Accepting this, we must have a measure of how valid our inferences based on old information will be as time elapses. DCFs provide this mechanism.

A third and final condition we put forth for using DCFs is that the expert system be applied to an environment where we are willing to "live with uncertainty." Specifically, the application must be such that we can accept the following proposition: although most decisions made will be profitable, some decisions may have negative effects simply because they are based on uncertain knowledge. For instance, in environments where decisions can make a difference between life and death of humans, we may not be willing to accept any uncertainty in knowledge. In such environments, we would require constant updating of information rather than using old information for inference. Since uncertainty is intolerable, having a method for providing time-varying degrees of confidence to knowledge based on aging information does not make sense. Fortunately, most environments though do not fall in this "matter of life and death" category.

4. Efficiency Gains Derived from Using DCFs

We hypothesize that there is great potential for gain in efficiency in the operation of expert systems using DCFs. This gain is derived, first of all, from making better decisions. An expert system incorporating DCFs takes into account not only the fact that information may be uncertain when it is learned, but also that uncertainty of inferences based on aging information increases as the information ages.

A second and more subtle reason for efficiency gain is due to the potentially economical use of all resources involved when learning new information. These can be resources at the source of the information (e.g., for preprocessing of information to be distributed), the actual resources supporting communication (e.g., networks and intermediate nodes), or resources used by a receiver (e.g., for postprocessing of received information). We illustrate this through an example.

Consider two entities, a machine which implements some function, and a controller, which controls the machine. This division of labor is necessary because the machine lacks processing power to control itself effectively. The controller monitors the machine and sends commands to it based on what it believes is the current state of the machine. The controller is assumed to make the optimal controlling decisions if it knows accurately the current state of the machine.

One way to determine the state of the machine is for the machine to periodically send updates to the controller. Each time an update is transmitted, resources as described above are used. Therefore, it is worth minimizing the frequency of updating. It would be highly desirable to find a way of sending updates only when necessary.

A reasonable approach is to send updates when the potential cost of making bad decisions (i.e., decisions that have undesirable effects), because they are based on old information, becomes greater than the known cost of simply sending an update. DCFs measure the uncertainty of knowledge as a function of the age of the information used to establish that knowledge. Consequently, there is some threshold of intolerable uncertainty of knowledge, where the potential cost of a decision based on that knowledge becomes greater than the cost of an update. Going beyond this threshold would trigger the need for an update.

Of course, this assumes the DCF is a measure proportional to the validity of aging knowledge. What if knowledge suddenly becomes invalid, but that its DCF still produces a high measure of confidence. Since the machine sends all updates, the machine would know that the controller has information that does not reflect current reality. It may also know the DCF used to model the decaying confidence of the controller's aging knowledge. Then, the machine can certainly recognize when the controller is highly confident of knowledge which is actually wrong. The machine would then take the initiative to send updated information.

Finally, we point out that if the mechanism which continually decays confidence measures according to the DCF is a built-in function of the expert system, we gain efficiency over systems which implicitly decay confidence measures using rules programmed by the user. In our expert system implementation, we associate with each item of knowledge, a DCF to measure confidence as time elapses, a pointer to a procedure, and a confidence measure threshold. When the DCF produces a value which goes below the threshold, the procedure given by the pointer is executed. We use this procedure to cause an update to either be requested or sent, depending on whether the role of the expert system is that of a controller or machine as in the example.

5. Summary

We have described a mechanism for taking into account decaying confidence of aging knowledge in expert systems. This general mechanism, called a Decaying Confidence Function, is applicable to all confidence measures. We put forward three conditions which justify the use of DCFs:

- the validity of learned information becomes uncertain as time elapses because the information describes time-varying entities;
- communication of information is costly, making it desirable to approximate current reality using old information;
- making bad decisions once in a while is acceptable.

We hypothesize that there is great potential for gain in efficiency in the operation of expert systems using DCFs. This gain is derived, first of all, from making better decisions, since an expert system incorporating DCFs takes into account not only that information may be uncertain when it is learned, but also that uncertainty of knowledge increases as the information it was based on gets older.

A second gain is derived from the potentially frugal use of all resources involved when new information is communicated. We can use DCFs as a guide to when it is necessary to communicate new information.

Finally, by making DCFs a built-in mechanism of an expert system, efficiency is gained over systems which implicitly implement the concept of increasing uncertainty as knowledge ages, through rules programmed by the user.

6. Acknowledgements

I would like to thank my advisor Professor Domenico Ferrari for his guidance, ideas, and valuable advice. Without his help and constant encouragement, this work would not be possible. Also, I wish to thank Professor David Anderson who originally suggested to me the possibility of confidence measures which decay as a function of time. Finally, I am very grateful to Dr. Ron Ashany for reading and making suggestions and corrections on earlier versions of this paper.

7. References

1. Shortliffe, E.H., *Computer-based medical consultation: MYCIN*, American Elsevier, New York, 1976.

2. Pearl, J., On Evidential Reasoning in a Hierarchy of Hypothesis, *Artificial Intelligence Journal*, 28:1 (1986) pp. 9-16.

3. Shortliffe, E.H. and Buchanan, B.G., A model of inexact reasoning in medicine, *Mathematical Biosciences*, 23 (1975) pp. 351-379.

4. Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press, 1976.

5. Zadeh, L.A., The role of fuzzy logic in the management of uncertainty in expert systems, *Fuzzy Sets and Systems*, 11 (1983), pp. 199-227.

6. Gaines, B.R., Fuzzy and probability uncertainty logics, *Information and Control*, 38 (1978), pp. 154-169.

7. Steeb, R., D.J. McArthur, S.J. Cammarata, S. Narain, and W.D. Giarla, Distributed Problem Solving for Air Fleet Control: Framework and Implementation, in *Expert Systems: Techniques, Tools and Applications*, P. Klahr and D.A. Waterman, eds., Addison-Wesley, Reading, MA, 1986, pp. 391-432.

8. Barstow, D.R., et al., Languages and Tools for Knowledge Engineering, in *Building Expert Systems*, F. Hayes-Roth, D.A. Waterman, D.B. Lenat, eds., Addison-Wesley, Reading, MA, 1983, p. 288.

9. Pasquale, J, "Using Expert Systems to Manage Distributed Computer Systems," *UCB Computer Science Tech. Report*, 1987, also submitted for publication.