

AN UNDERLYING MODEL FOR DEFEAT MECHANISMS

Brent Heeringa
Paul Cohen

Experimental Knowledge Systems Laboratory
Department of Computer Science
University of Massachusetts, Amherst
Amherst, MA 01003, USA

ABSTRACT

Defeat mechanisms are strategies for achieving victory over an opponent. Although defeat mechanisms often rely on influencing the opponent psychologically and emotionally, most simulations of warfare do not model these “soft” factors, they model only victory by attrition. To create more accurate, adaptable, and believable systems, we must be able to model a variety of defeat mechanisms. We propose a model where parameters and attributes that affect emotional and physical fatigue are combined to produce an overall measure of fatigue called effective fatigue. Effective fatigue, along with an agent’s state, is combined by a defeat model to produce probabilities of surrender. We create warfare scenarios involving catastrophe and surprise, and then examine the model’s behavior under these scenarios. We conclude with a discussion of how the model is related to our own Capture the Flag wargaming system.

1 INTRODUCTION

Frequently, the goal of military action involves making one’s opponent capitulate, so the study of military action includes *defeat mechanisms*, or strategies for achieving capitulation. Defeat mechanisms include the element of surprise, catastrophe, as well as victory by attrition (Clausewitz 1976) (Tzu 1963). Surprise means catching an agent off-guard both psychologically and physically, catastrophe means inflicting significant damage in a short interval, and victory by attrition involves persistent damage until an agent surrenders or is destroyed. One view of defeat is that the warrior has a limited supply of psychological and physical resources, and that defeat occurs when these resources are used up. Courage, for example, is considered by Lord Moran to be a “a moral quality” that is spent over time (Moran 1945). While grinding attrition undoubtedly depletes a warrior’s psychological resources, other defeat mechanisms might bring about capitulation more




Figure 1: The Capture the Flag Wargame Simulator

quickly. However, it is difficult to empirically evaluate various defeat mechanisms and combinations of defeat mechanisms, because modern wargaming systems model only victory by attrition (Zimm 1999). While military theorists design maneuvers explicitly to affect the psychological state of their opponents, they lack the simulation tools to evaluate these effects. A wargaming system that accurately models factors of fatigue, and their effect on an agent’s probability of surrender is more accurate, in a predictive and explanatory sense, than one that does not.

We have developed a wargame simulator called Capture the Flag (CtF). Using CtF (see Figure 1), we can predict and explain courses of action (COA) in war. We have recently added fatigue and defeat models to

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2000		2. REPORT TYPE		3. DATES COVERED 00-00-2000 to 00-00-2000	
4. TITLE AND SUBTITLE An Underlying Model for Defeat Mechanisms				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts, Department of Computer Science, Amherst, MA, 01002				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 2: The Abstract Fatigue Model

CtF thereby increasing the accuracy of our simulator, creating more adaptive behavior in planning for defeat, and allowing us to better explain battles and their outcomes.

There is little research in the area of modeling fatigue and defeat in military warfare, and much of it is inconclusive (Hudlicka and Billingsley 1999). In this paper, we propose abstract models for fatigue and defeat. We identify measurable parameters that affect physical and emotional fatigue, such as attrition and the proximity of opposing troops. Our fatigue model combines these parameters along with other attributes such as fear and courage to produce an overall measure of fatigue called effective fatigue. Our defeat model combines effective fatigue with an agent’s state to compute a probability of surrender.

2 MODELING FATIGUE

In our system, we do not try to simulate psychological or physiological processes in individual warriors, but instead we model the collective fatigue of a unit (e.g., a battalion) as a weighted sum of factors that *influence* fatigue. Fatigue is a function of its physical, emotional, and personal components. Physical fatigue can be thought of as a depletion of energy or mass, while emotional fatigue summarizes the effects of sensations such as fear, courage, aggression, and morale. Personal components are traits inherent to agents; for example, an agent’s warfare style may designate that it always fights to the finish or is quick to surrender.

Personal traits and factors influencing physical fatigue are often directly observable (e.g., warfare style, hours without sleep, attrition, length of current battle, etc.). In contrast, factors affecting emotional fatigue are difficult to measure directly.

Our fatigue model consists of parameters, which are directly measurable quantities, and attributes, which are not directly measurable. The values of parameters are supplied directly to the model via the person building it, or through values present in an external system, while attributes are variable and are influenced by other attributes and parameters in the model. Effective fatigue is a combination of parameter and attribute values. Later, in Section 3.1, we will show how effective fatigue is combined with information about an agent’s *state* to produce an overall probability of surrender.

2.1 The Abstract Fatigue Model

Figure 2 represents our abstract fatigue model. Formally, the model is a four-tuple $\mathcal{FM} = \langle P, A, F_e, \alpha \rangle$ where:

- $P = \langle p_0, p_1, \dots, p_{n-1} \rangle$ = a vector of parameters
- $A = \langle a_0, a_1, \dots, a_{m-1} \rangle$ = a vector of attributes
- F_e = effective fatigue
- $\alpha = \langle \alpha_0, \alpha_1, \dots, \alpha_{q-1} \rangle$ = a vector of influence arcs

Parameters and attributes are connected to effective fatigue through arcs, which represent influence. Each arc $\alpha_i = \langle n_0, n_1, \psi \rangle$ is a three-tuple consisting of a from-node (n_0), a to-node (n_1), and an influence function (ψ). The influence function ψ allows us to control the effect, or influence, parameters and attributes have on effective fatigue. We allow outward-pointing arcs (feed-forward arcs) from parameters and attributes to effective fatigue, but conversely, inward-pointing arcs (feed-back arcs) are directed at attributes only.

Zimm (1999) justifies the feedback arcs, noting:

- destruction causes panic and paralysis; and
- panic and paralysis facilitates destruction.

Moreover, since attribute values are not directly measurable, our modeling language provides means for calculating those values as combinations of measurable quantities (i.e., parameters). Detailed discussion and concrete examples will follow in Section 4.

3 THE DEFEAT MODEL

A defeat model contains a base probability of surrender, a set of states, rules for specifying when state transitions are made, and functions that specify how the current probability of surrender is computed based on the time spent in the current state.




Figure 3: A Possible INBATTLE Sub-state

Every defeat model has an initial base probability of surrender. This base probability is purely a function of effective fatigue. In addition, modelers may define other states. These additional states modify the initial probability to produce an agent's final probability of surrender.

The defeat model's estimation of an agent's probability of surrender is based heavily on an agent's current state in war. For example we may occupy a PREPAREDFORBATTLE state, that indicates we are currently not in, but prepared for, battle. Each state is comprised of sub-states. This removes the complexity of war, by decomposing situations into identifiable units. For example, the PREPAREDFORBATTLE state may be comprised of PREPARED and NOTINBATTLE sub-states. Collectively, sub-states describe an agent's current situation in war. Each state combines its sub-states to compute an agent's overall probability of surrender.

3.1 Defeat Model States

Formally, a state is a three-tuple $S_i = \langle \omega, \mathcal{C}_{S_i}, \lambda \rangle$ where:

- ω = a set of sub-states.
- $\mathcal{C}_{S_i}$ = a set of criteria for transitioning into state S_i .
- $\lambda_i = \lambda_i(\omega) =$ the probability of surrender for state S_i = a combination function over our set of sub-states ω .

Each sub-state ω_i is composed of a modifier function ρ and a set of criteria, $\mathcal{C}_{\omega_i}$, for state transition. For example, the INBATTLE sub-state in Figure 3 states that we are in battle when we were not in battle and suddenly incur damage, or when an opponent is 5 units




Figure 4: A Possible Combination of Sub-states

of distance away from us. The function ρ denotes how long we have occupied the sub-state ω_i . It is also used by the combination function λ_i to modify the overall probability of surrender. For example, in Figure 3, we see that the INBATTLE modifier increases the probability of surrender at the start of a battle, but over time, decreases its influence.

$\mathcal{C}_{S_i}$ is the union of transition criteria $\mathcal{C}_{\omega_i}$, for every sub-state ω_i .

λ_i is the combination function. It provides a means of computing an overall probability of surrender based on our set of sub-states ω and the base probability of surrender $\mathcal{B}$.

For example, Figure 4 illustrates one possible combination function. Given n sub-states, we arbitrarily order them $\omega_1, \dots, \omega_n$. First, ω_1 calculates its modifier value by computing, $\lambda_1(\rho_1(t_1), \mathcal{B})$, where t_1 is the time we have occupied state ω_1 and $\mathcal{B}$ is our base probability of surrender. Next ω_2 computes its modifier value based on ρ_2 and λ_1 . We continue this process until we reach sub-state ω_n , where λ_n denotes the agent's final probability of surrender.

4 MODELING CATASTROPHE AND SURPRISE

In this section, we create example fatigue and defeat models and view the effects of catastrophe and surprise scenarios on the model. We also analyze and examine the overall behavior of the fatigue and defeat models.

4.1 An Example Fatigue Model Instance

Each warfare system is different. To make the model accurate, the designer of the wargame system must answer




Figure 5: An Instance of the Fatigue Model

questions such as “What levels of fatigue are high?” and “How much damage is usually incurred during a given period of time?”

Figure 5 represents an instance of our fatigue model $\mathcal{FM}$. In our warfare system, 450 units of effective fatigue is high and, in battle, 10 units of damage per tick is typical. That is, when an agent’s effective fatigue level reaches 450 units of damage, we should start seeing significant increases in its respective probability of surrender. It is worth noting that each agent in our system is representative of a battalion or brigade. Hence, for our purposes, the effect of each catastrophe and surprise scenario is not measured on an individual level, but at a higher resolution.

$\mathcal{FM}$ consists of four attributes and two parameters (see Table 1). Each attribute a_i and parameter p_i has value in the open interval $(0, 1)$, except for *attrition* which has value in the open interval $(0, +\infty)$.

Courage represents an agent’s spirit and tenacity. We define $a_{courage} = 0$ as feeling extremely courageous and $a_{courage} = 1$ as a total lack of courage, or even a state of frenzy. Moran (1945) suggests courage may help in battle, thus our arc is weighted on the interval $(-10, 10)$. Note that negative weight values allow $a_{courage}$ to lower effective fatigue.

Health represents an agent’s level of sickness. We define $a_{health} = 0$ to be completely healthy, and $a_{health} = 1$ to be deathly ill. We model health as an attribute because an agent’s level of sickness is difficult to mea-

Table 1: Our Fatigue Model Attributes and Parameters

<i>name</i>	<i>type</i>	<i>symbol</i>
courage	attribute	$a_{courage}$
fear	attribute	a_{fear}
health	attribute	a_{health}
morale	attribute	a_{morale}
warfare style	parameter	p_{style}
attrition	parameter	$p_{attrition}$

sure directly. Note that health is different from attrition, but can both influence and be influenced by attrition indirectly, through effective fatigue. Our feedforward arc uses a simple function $f(x) = 30x$ to account for the effect of health on effective fatigue. Essentially then, we can view our feedforward arc as having weight $30 * a_{health}$.

Morale represents an agent’s level of confidence, enthusiasm and sense of purpose. We say $a_{morale} = 0$ means morale is high, and $a_{morale} = 1$ indicates morale is low. The feedforward arc has weight $(60 * a_{health}) - 30$. That is, a_{morale} is mapped into the open interval $(-30, 30)$.

Fear represents an agent’s level of trepidation. We say a unit is feeling no fear when $a_{fear} = 0$ and filled with fear when $a_{fear} = 1$. Fear has an associated arc weight of 40 units.

Warfare style characterizes a bias in battle style inherent to an agent or group of agents. We say that $p_{style} = 0$ means *under no circumstance* will a unit surrender, and that $p_{style} = 1$ means *under most circumstances* a unit will surrender. To model this correctly, we attach some parabolic function (say $f(x) = 5000(x - .5)^2$), that allows us to create overwhelming influences on effective fatigue. For example, if an agent’s warfare style is to surrender quickly and easily, our arc will produce negative weights that consume all other influences. That is, the arc weight value is so low, that all other attribute and parameter effects on effective fatigue are rendered meaningless.

Attrition alone comprises this model’s representation of physical fatigue. An agent’s attrition level (i.e. combat attrition) is provided directly from the wargame system. That is, after every tick t , the system updates $p_{attrition}$ (via the incoming arc) to reflect an agent’s current damage level.

While we want to propagate the actual value of attrition forward to effective fatigue, we are also interested in its rate of change. The rate of change of attrition helps indicate when significant changes in battle occur.




Figure 6: The Base Probability of Surrender

For example, a high rate of change in attrition may indicate a catastrophe, while a sharp decrease may indicate medical relief.

We map rates of change into the interval $(-40, 40)$. When rates of change are positive, a higher value from our interval is added to attrition. When rates of change are negative, a lower value from our interval is subtracted from attrition. No rate of change in attrition maps to 0 in the interval.

Effective fatigue has four arcs, providing feedback to the morale, fear, health, and courage attributes. Each arc has an associated function which first finds the discrete derivative of effective fatigue over one time unit, and maps that derivative to a multiplier m , where $0 < m < 2$. Each value v_i in attribute a_i is then set to $m * a_i$. The discrete derivative allows us to model significant increases or decreases in effective fatigue, and in turn, provide corresponding feedback to the attributes. We don't model rate of change on arcs from attributes to effective fatigue because the feedback mechanism indirectly provides such a measure.

It is also worth noting that the effective fatigue value for time t is computed by simply summing the values returned by each of its feedback arcs.

4.2 An Example Defeat Model Instance

Our defeat model $\mathcal{DM}$ has a simple base probability of surrender function based on the exponential distribution:

$$\mathcal{B}(F_e) = \alpha \lambda e^{-\lambda(\beta - F_e)} = 40 * .0125e^{-.0125(450 - F_e)} \quad (1)$$

Equation 1 is depicted in Figure 6. β determines where the mean ($\alpha * \lambda$) of the distribution will occur. Changing the α and λ values affects the convexness of the exponential arc. In the beginning of Section 4 we

noted that 450 units of effective fatigue is significant in our system. Because of this, we chose $\beta = 450$.

Our defeat model, $\mathcal{DM}$, contains four states:

- PREPARED/INBATTLE
- UNPREPARED/NOTINBATTLE
- PREPARED/NOTINBATTLE
- UNPREPARED/INBATTLE

composed of four sub-states:

- PREPARED
- UNPREPARED
- INBATTLE
- NOTINBATTLE

PREPARED denotes the state of being prepared for battle. UNPREPARED denotes the state of being unprepared for battle. Both states use their respective modifier and combination functions to respectively decrease and increase the input probability distribution by a some percentage.

INBATTLE denotes the state of currently being in battle and NOTINBATTLE denotes the state of currently not being in battle. The BATTLE modifier use a variant of the exponential distribution to produce higher modifier values at the beginning of battle. The NOTINBATTLE modifier uses the identity function to leave the incoming probability of surrender unchanged.

4.3 Catastrophe and Surprise

Catastrophe and surprise are two unique defeat mechanisms used in warfare. Catastrophe relies heavily upon inflicting massive physical damage on an agent in a relatively short period of time. In contrast, surprise is more psychological in nature. It seeks to catch an agent physically and emotionally off-guard in order to promote surrender quickly and effectively.

We created a wargame system that models an agent's level of attrition over time. Using our system, we simulated normal combat, catastrophe and surprise scenarios. These scenarios were created through changes in an agent's level of attrition over time. We created instances of the example fatigue and defeat models given in Sections 4.1 and 4.2 and, after each tick of the simulation, provided the agent's current level of attrition to the model. We then recorded the agent's probability of surrender for each respective time frame.




Figure 7: Attrition Level for Catastrophe Scenarios




Figure 8: Probability of Surrender for Catastrophe Scenarios

4.3.1 Catastrophe

Figure 7 depicts three scenarios in which catastrophe occurs at the beginning ($C_{beginning}$), middle (C_{middle}), and end (C_{end}) of the battle respectively. The x-axis denotes time, which in this case, corresponds to ticks of the simulator. The y-axis corresponds to an agent's probability of surrender at any given tick. These catastrophe scenarios only transition between the PREPARED/NOTINBATTLE and PREPARED/INBATTLE states.

Figure 8 depicts the respective probabilities of surrender for each catastrophe scenario. We see that in each case, when a catastrophe occurs, the probability of surrender significantly increases.

$C_{beginning}$ contains a sharp increase in surrender because the catastrophe is significant and it occurs at the beginning of a battle when the probability of surren-

der is higher. As soon as the battle begins though, it ends, thus the sharp decrease in probability of surrender. This sudden drop is probably too dramatic. It may be useful to introduce an intermediate state BATTLEOVER between INBATTLE and NOTINBATTLE that prolongs the effects of a completed battle over some time period. Later, at around tick 30, we see another small catastrophe, and correspondingly, an increase in the probability of surrender. The increase in probability of surrender is fairly high considering only the small catastrophe, however the cataclysm increased our damage to significant levels.

C_{middle} contains a medium-grade catastrophe during the middle of a battle. The calamity occurs after a steady, constant increase in attrition, and hence the corresponding increase in probability of surrender is also somewhat mid-grade. Notice that the probability of surrender, up to the point of the disaster, is extremely small and constant. This behavior seems fitting as the rate of change in attrition levels is constant, and moreover, those attrition levels are relatively low.

C_{end} contains three small catastrophes in succession during the final 20 ticks. It appears the first two catastrophes only slightly increase the probability of surrender. This is justified by a number of factors. First, the catastrophes occurred at times when attrition was increasing. Next, the attrition levels were not at significant levels to propagate higher surrender probabilities, and finally, the catastrophes were fairly insignificant. The final catastrophe increases attrition to a significant level and hence, the dramatic increase in the corresponding probability of surrender. One interesting behavior is the sharp lowering of the surrender probability near the end of the battle. This sudden drop seems wrong. A new intermediate state after BATTLE may improve the model's behavior.

There has been a significant amount of work done in depicting catastrophe through smooth functions. Bifurcation theory attempts to fit a smooth function along with a constant factor to a time-series. Through small changes in the constant factor, discontinuities occur in the smooth function, hence reflecting catastrophe (Casti 1989). Thus, bifurcation theory may help indicate if our model behaves correctly under catastrophe.

Using the probability of surrender time-series C_{end} , we can easily fit the smooth function $\alpha * e^x$ to the curve using extremely low constant values (on the order of $6.0 * 10^{-25}$) for α (similar curve fittings are possible for $C_{beginning}$ and C_{middle}). Of course, the sudden drops in the probability of surrender would indicate a positive catastrophic event. Clearly, in our scenarios, this is not the case.




Figure 9: Attrition Level for Surprise Scenarios




Figure 10: Probability of Surrender for Surprise Scenario

4.3.2 Surprise

Figure 9 represents the attrition level for the surprise scenarios. The scenarios are identical, hence, only the one curve in Figure 9, except in one scenario, the agent is always prepared for battle, while in the other, the agent is always unprepared. This scenario contains steady increases in attrition over time, not unlike a typical battle.

Figure 10 represents the probability of surrender associated with the given scenarios. Note the minor differences in the probability of surrender during the first state change. This minor escalation is due to low levels of attrition. In contrast though, as attrition begins to rise at tick 30 following a state change, we see a sudden large influence in the associated probability of surrender. Due to a high and consistently increasing attrition level, this increased probability of surrender is easily justified.

5 DISCUSSION

We tested the model on catastrophe and surprise scenarios. The catastrophe scenarios reflected a significant increase in an agent's attrition over a short period of time. The surprise scenarios used identical attrition values, but designated different states for the agent. Under both test scenarios, the model's behavior was fairly believable. The one caveat occurred soon after catastrophes occurred, with sudden, dramatic drops in the probability of surrender. That is, the model behaves well when catastrophe first occurs, but is a bit more unpredictable and sporadic after such calamities. Finding the proper balance of states and probability modifiers is certainly an area worth further investigation.

6 FUTURE WORK

For some time now, the Experimental Knowledge Systems Laboratory has been developing tools for simulating physics abstractly, for planning in dynamic, real-time environments, and for hierarchical agent control (Atkin et al. 1998) (Atkin and Cohen 1998) (Atkin, Westbrook, and Cohen 2000). This work has led to the creation of a warfare simulator called Capture the Flag. Capture the Flag, like other warfare simulators, uses a lanchester-based attrition model and suffers from the "attrition paradigm." We have recently incorporated the fatigue and defeat models into Capture the Flag. It will now be possible to further explore the dynamics of battle, comparing warfare simulations that incorporate defeat mechanisms with those that do not.

ACKNOWLEDGMENTS

We would like to thank General Charlie Otstott and the anonymous reviewers for their helpful and insightful comments.

This research is supported by DARPA/USAF under contract numbers N66001-96-C-8504, F30602-97-1-0289, and F30602-95-1-0021. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements either expressed or implied, of the Defense Advanced Research Projects Agency/Air Force Materiel Command or the U.S. Government.

REFERENCES

- Atkin, Marc S. and Paul R. Cohen. 1998. "Physical Planning and Dynamics." *Working Notes of*

the AAAI Fall Symposium on Distributed Continual Planning. 4–9.

Atkin, Marc S., David L. Westbrook, and Paul R. Cohen. 2000. “Domain-General Simulation and Planning with Physical Schemas.” To appear in the Proceedings of the 2000 Winter Simulation Conference.

Atkin, Marc S., David L. Westbrook, Paul R. Cohen, and Gregory D. Jorstad. 1998. “AFS and HAC: Domain-General Agent Simulation and Control.” *Working Notes of the Workshop on Software Tools for Developing Agents, AAAI-98.* 89–95.

Casti, John L. 1989. *Alternate Realities.* John Wiley and Sons.

Clausewitz, Carl Von. 1976. *On War.* Princeton University Press.

Hudlicka, Eva and John Billingsley. 1999. “Representing Behavior Moderators in Military Human Performance Models.” *8th Conference on Computer Generated Forces and Behavioral Representation.*

Moran, Lord. 1945. *Anatomy of Courage.* Avery Publishing Group Inc.

Tzu, Sun. 1963. *The Art of War.* Oxford University Press.

Zimm, Alan D. 1999. “Modeling Maneuver Warfare: Incorporating Human Factors and Decisionmaking in Combat Modeling.” *Proceedings of the 8th Conference on Computer Generated Forces and Behavioral Representation.*

in 1993 as a Fellow of the American Association for Artificial Intelligence, and served as Councillor of that organization, 1991–1994. His research concerns the design principles for intelligent agents and the acquisition of conceptual structures. His email and web addresses are <cohen@cs.umass.edu> and <eksl.cs.umass.edu/~cohen>.

AUTHOR BIOGRAPHIES

BRENT HEERINGA received his B.A. in Computer Science and Mathematics at the University of Minnesota, Morris. Currently, he is working toward his Ph.D. at the University of Massachusetts, Amherst in the Experimental Knowledge Systems Laboratory under the direction of Paul R. Cohen. His research interests include unsupervised learning, simulation, formal methods and game theory. In his spare time he enjoys independent music and film. His email and web addresses are <heeringa@cs.umass.edu> and <www.cs.umass.edu/~heeringa>.

PAUL R. COHEN received his Ph.D. from Stanford University in 1983. He is a Professor of Computer Science at the University of Massachusetts, and Director of the Experimental Knowledge Systems Laboratory. He edited the “Handbook of Artificial Intelligence,” Volumes III and IV with Edward Feigenbaum and Avron Barr. Cohen was elected