

(2)

AD-A258 322



Adapting the Contact Migration Policy Algorithm to Dynamic Planning in Phoenix

Paul R. Cohen, *Principal Investigator*
David M. Hart, *Lab Manager*

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

cohen@cs.umass.edu
dhart@cs.umass.edu
413 545 3638

DTIC
ELECTE
DEC 7 1992
S C D

Final Report

Contract N00014-91-J-1750
Period 5/1/91 - 4/30/92

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

92-30915



24pgs

92

4

Table of Contents

Numerical Productivity Measures	1
Brief Summary of Technical Results	2
Detailed Summary of Technical Results	4
Publications, Presentations, Reports, Awards/Honors	13
Transitions and DoD Interactions	15
References	17
Appendix A: Textbook Prospectus	20

DTIC QUALITY INSPECTED 3

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

1 Numerical Productivity Measures

Refereed papers: 2

Unrefereed reports and articles: 1

Books or parts thereof: 1

Patents filed but not yet granted: 0

Patents granted: 0

Invited presentations: 7

Contributed presentations: 2

Honors, including conference committees: 9

Paul Cohen:

Councilor, American Association for Artificial Intelligence, 1991-1994.

Co-Chair, Symposium Committee, American Association for Artificial Intelligence, 1992-1993.

Chair, Tutorial Committee, American Association for Artificial Intelligence, 1992-1993.

Assistant to the Chair, Program Committee, American Association for Artificial Intelligence, 1992-1993.

Member, Information and Science Technology Advisory Group on Simulation, Institute for Defense Analysis, 1992-1993.

Co-Chair, Tutorial and Symposium Committees, American Association for Artificial Intelligence, 1991-1992.

Program Committee, *Sixth International Symposium on Methodologies for Intelligent Systems* (ISMIS'91). Charlotte, NC, October 1991.

Organizing Committee, *AAAI Workshop on Intelligent Real-Time Problem Solving*. Anaheim, CA, July 1991.

Chairman, *NSF/DARPA Workshop on AI Methodology*. University of Massachusetts, June 1991.

Promotions: 0

Graduate students supported at least 25% time: 4

2 Brief Summary of Technical Results

2.1 Applying the Contact Migration Policy to Phoenix

The *contact migration policy* (Gruen 1991) is a reactive approach to constraining (grasping) a rigid body that incrementally constructs a stationary process model of the body in order to refine the grasp. The algorithm employed has many of the characteristics of an anytime algorithm (Dean and Boddy 1988), and is therefore an attractive candidate for real-time applications such as positioning bulldozers in the Phoenix fire-fighting simulation. However, we have found that several of the assumptions underlying this algorithm do not hold in the Phoenix environment, so that such a purely reactive control mechanism is inadequate in Phoenix (see Section 3.1). Based on these findings, and on a broader survey of the application of reactivity in robotic systems, we argue for the mixture of reactive and deliberative (or predictive) mechanisms to coordinate the fire-fighting activities of Phoenix agents. The survey on reactivity has led to the development of a graduate seminar on reactive robotic systems being offered this fall.

2.2 Resolving the Deliberative/Reactive Dilemma

We pursued this argument – that the control architecture in Phoenix should combine deliberative and reactive control – by following several related lines of research. In the first, we explored the effects of placing deliberative and reactive controllers on separate processors within the same architecture (as in Phoenix). We derived a simple mathematical model of the interaction of two tasks running concurrently on a uniprocessor. We modeled the rate at which two periodic tasks would interrupt each other. From the model, we argued that placing reflexive control and cognitive control on different processors is not justified by the different time-scales over which they act, but is justified by the extent to which the higher priority task dominates the uniprocessor.

The second line of research addressed the question, *when should an agent deliberate and when should it react?* We have attempted to answer this question by establishing a set of principles and criteria that govern this decision. To evaluate these guidelines we implemented an artificial world for path planning (navigation). This artificial world consists of linear obstacles of random length and orientation. These obstacles, called "chaff" hence the world was called "Chaffworld," can be avoided more efficiently by planning a path around them (deliberating) rather than bumping into them (reacting), but this efficiency is bought at the cost of that planning.

Based on experiments run in Chaffworld (see Section 3.2), we claim that the design of an agent must reflect the task it undertakes and the environment it experiences. In some environments, the agent will be reactive, while in others it will be deliberative. In environments that exhibit qualities of both reactive and deliberative environments, the agent design may well be a "two-loop" design (as in Phoenix). We show that we can design the right kind of agent for its environment by knowing the cost/benefit

tradeoff in the task, the level of abstraction of the representation used in deliberation, and the entropy of the environment.

2.3 Follow-up on Previous ONR-sponsored Research

We have pursued several follow-up activities to our previous ONR funded research and take this opportunity to report on those. We specifically report continued progress in two areas:

A Textbook for Empirical Methods in Artificial Intelligence. We have been preparing a textbook based on the methodology we have developed for conducting research in AI, and are using that text for a graduate seminar this fall. A prospectus of the textbook appears in Appendix A.

An Automated Model-Building Assistant. We are developing an *automated model-building assistant* capable of extracting causal relationships from data and constructing analytical models to guide the design of autonomous agents. The long-term goal of this research is to embed this system in an autonomous agent architecture, *creating agents capable of integrating perceptions of their task environments into conceptual models that support reasoning and planning.*

3 Detailed Summary of Technical Results

3.1 Applying the Contact Migration Policy to Phoenix

3.1.1 Specification of metrics of performance.

In the grasping problem discussed in (Gruppen 1991, Gruppen & Weiss 1991), metrics of performance that express context dependent subgoals of the task (grasp formation) are:

- *Null space error.* This term expresses the error of the contact system with respect to the task and is unimodal.
- *Manipulability of the hand.* A convex composition of the manipulability of the fingers, this term expresses the kinematic quality of the grasp configuration.

Admissible and predictable controllers for grasp optimization are constructed from the above mentioned formal metrics. Cooperative reactivity emerges from the constraints imposed by the manipulability of the hand on migrating contacts.

In Phoenix, no suitable analogs to these formal metrics of performance exist currently, nor were we able to construct any. Predictable and admissible controllers constructed from formal metric specifications are prerequisites to cooperative reactivity using the contact migration policy.

3.1.2 Process Modeling.

In the grasping problem, uncertainty arises due to incomplete information about the object (to be grasped). Gruppen's strategy for grasp formation refines an incomplete process model from local information obtained at individual contact sites. This information is used to perform reactive adjustments, to optimize the grasp locally. However, the object to be grasped is *rigid* and *static*. The kinematics of the hand are *fixed*. Further, contact sites provide sufficient process sampling density. These relative advantages simplify the construction of a stationary process model.

In contrast, parameters of fire dynamics change continuously, in highly nonlinear ways as fire consumes forest, traversing varying terrain and subject to changing wind velocity. This entails the construction of a non-stationary and dynamic model. This a harder problem and requires a much more densely sampled process than that required in grasping a rigid body. Phoenix observers (i.e., the bulldozers fighting the fire) are however inherently sparse. We have determined that the process modeling problem in Phoenix requires a different approach than that used in contact migration – one sufficiently different that the contact migration algorithm is no longer applicable.

3.1.3 Reactivity in Phoenix

The Phoenix environment has been characterized as unpredictable, dynamic and complex (Cohen et al. 1989). Reactive schemes have been employed in unpredictable environments, such as in mobile robot navigation discussed in (Arkin 1987, Arkin

1988), (Payton 1986, Payton 1990) and control theoretic approaches have been employed to construct models for reactive interaction between the agent and the environment (Khatib 1985). However, purely reactive schemes have been observed to fail in complex environments (Ravela 1992), whereas global planning approaches seem to be more suitable (Chrisman 1991). Indeed, planning has been integrated with reactivity to tackle complex, uncertain environments; mechanisms for integrating planning and reactivity have been discussed in (Payton 1990; Kaebbling 1989; Anderson, Hart & Cohen 1991; see below).

An understanding of the modeling problem in Phoenix suggests that, though uncertainties may warrant a reactive component, process modeling must however evolve from predictive capabilities (expressed as contingency anticipation), domain knowledge and map-based reasoning (expressed as plans or expectations), and incremental evidence assimilated from reactive behavior. Chrisman (1991) describes a similar characterization. Further, the specification of tasks and plans, the construction of consistent world models and models of interaction between the agent and the environment are all issues that form the basis of a design methodology for cooperative reactivity.

3.1.4 Conclusions

Simultaneous effects of, limited observability, complex, non-stationary, nonlinear process dynamics and the presence sparse observers inhibit cooperative reactivity (as defined for the contact migration policy) in Phoenix. The Phoenix domain represents a dramatically more difficult problem for the behavioral design paradigm than that for which the contact migration policy was developed. Significant developments in modeling coupled with knowledge/map based reasoning and heuristic models of weather variability are prerequisites to developing cooperative reactivity.

3.1.5 Survey and Graduate Seminar on Reactivity

Preliminary to the above investigations we conducted a broad survey of the application of reactivity in robotic systems (Ravela 1992). This survey provided a context for these investigations and is also the basis for a recently developed graduate seminar entitled "Reactive Robotic Systems" being offered in the fall of 1992 for the first time. The seminar description reads:

This seminar will focus on the specifications for and application of reactive subsystems to integrated robotic systems. We will review recent work in formal specifications for reactive systems which yield provably convergent controllers. The focus will be on sensor-based navigation functions for mobile platforms and on the control composition problem for high dimensional systems with excess degrees of freedom. The mobile platform study will deal with issues of dimensionality, and the fusion of information from vision, sonar and touch into a consistent environmental model. We will study the application of new techniques for reactive re-planning in this context. The redundant systems study will focus on issues related to building controllers for the Utah/MIT dextrous hand. The emphasis here is the decomposition of high dimensional state spaces into tractable subspaces and the dynamic, on-line re-composition of controllers mediated by sensory information and the task. We will study methods for

building structure into the re-composition problem using domain knowledge, and learning composition policies.

3.2 Resolving the Deliberative/Reactive Dilemma

One of the central dilemmas in AI in the past several years has been deliberative planning versus reactive planning. The former is what AI has called planning since the time of STRIPS, while the latter is a reaction to the computational complexity of formal planning and emphasizes algorithms that simply react to the current situation.

From the beginning of the Phoenix project, we saw a need for both kinds of capability: the foresight of a deliberative planner in order to coordinate the resources of multiple firefighting field agents, and the speed of a reactive planner in order to survive sudden changes in the direction and location of the fire. Consequently, we designed what we have called a "two-loop" architecture (see Figure 1) where the outer sense-act loop is slow enough to accommodate deliberative planning while the inner sense-act loop is quick enough to react to a changing environment.

In our initial investigations of this design's efficacy, we derived a simple mathematical model of the interaction of two tasks running concurrently on a uniprocessor (Anderson, Hart & Cohen 1991). Specifically, we modeled the rate at which two periodic tasks would interrupt each other. The predictions of this model were supported by simulation experiments. From the model, we argued that placing reflexive control and cognitive control on different processors is not justified by the different time-scales over which they act, but *is* justified by the extent to which the higher priority task dominates the uniprocessor.

The research question we've been investigating since then is essentially: *when should the agent deliberate and when should it react?* What are the dimensions of that decision? We believe the answer rests on the following principles and criteria:

- Deliberation, especially "Forward Search," is computationally expensive
- Deliberation is sometimes useful or even necessary
- Its use is dictated by:
 - Computational Cost vs Execution Benefit
 - Abstraction Level
 - Uncertainty
 - * Intrinsic unpredictability of the world
 - * How much precision does the agent demand?




Figure 1. The upper picture views agent behavior very abstractly as a mapping from perceptions to articulator trajectories. The lower picture depicts our modularization of that abstract mapping. The interface between the modules consists of trajectories of subgoals (subgoals of the overall behavior). Adding perceptual input to the interface allows for a quick feedback loop, thereby allowing reactivity.

To attack the question of cost and benefit, we implemented an artificial world for path planning (navigation). This artificial world consists of linear obstacles of random length and orientation. These obstacles, called "chaff" hence the world was called "Chaffworld," can be avoided more efficiently by planning a path around them (*deliberating*) rather than bumping into them (*reacting*), but this efficiency is bought at the cost of that planning. An idealized chaff encounter is depicted in Figure 2. A typical chaffworld is shown in Figure 3.

We used an off-the-shelf A* search algorithm to do the planning, rather than a special purpose algorithm, so that the results will not be specific to navigation, but be more general. The mathematical analysis of prototypical encounters (Figure 2) suggested that when chaff is very sparse, planning is probably not worthwhile, but as the chaff density increases, the advantage of planning increases, so that a break-even point may be reached. This break-even point is the thinking ratio, c (measured as the ratio of the map area it can search per unit time to the linear distance it can travel in unit time), of an agent that is indifferent between deliberating and reacting. An agent with a higher ratio would prefer to deliberate, while an agent with a lower ratio would

prefer to just start moving. Thus, we have hard, objective reasons to deliberate or react.




Figure 2. An agent encounters an obstacle. It can either plan its way around ($\overline{AC}$) or bump into it and follow it around ($\overline{ABC}$).




Figure 3. The scene on the right is of a typical, randomly generated Chaffworld. The scene on the left is the agent's map of the Chaffworld. In some experiments, this map was identical to the world, while in others, shown here, short pieces of chaff are not represented, to test the hypothesis that abstraction of this sort will improve the efficiency of planning.

The density of chaff was controlled by two parameters used to generate the world: λ the mean of the exponential distribution of chaff lengths or the constant length of all chaff, depending on the experiment; and ρ , the rate of the Poisson process used to scatter the chaff across the world. The length of the chaff, λ , is measured in meters, while the rate, ρ , is measured in pieces of chaff per square kilometer. The experimental and theoretical results are used to compute c_{min} – the break-even thinking ratio, measured in square meters per meter. The numerical values for reasonable values of λ and ρ are presented in Table 1. The experimental results are mixed: they do suggest tradeoffs and increasing advantage for planning, but the empirical break-even point is quite different from its theoretical value. Further experiments are being designed to explain this discrepancy. One hypothesis is that the experimental situation (see Figure 3) cannot be adequately modeled as a random iteration of the idealized situation (see Figure 2), because a locally optimal path in the idealized situation is found, while the A* search solves the experimental situation for the globally optimal path. Therefore, the experiment should be redesigned so that the experimental situation more closely corresponds to the locally optimal solutions.

Theoretical Predictions of c_{min} .

ρ	λ			
	1000	2000	3000	4000
.1	12,570	7,260	7,290	6,410
.2	6,740	4,490	4,420	
.3	4,900			

Experimental Values for c_{min} .

ρ	λ			
	1000	2000	3000	4000
.1	224.1	14.3	19.6	8.4
.2	47.3	4.5	4.3	
.3	11.7			

Table 1. In this experiment and theoretical calculation, chaff was assumed to be distributed with constant length λ and density ρ . The theoretical and experimental results are qualitatively in agreement: smaller λ or smaller ρ produce a smaller c_{min} . Obviously, they strongly disagree on the actual value of c_{min} .

Chaffworld also partially addresses the question of abstraction level. The abstraction level is the primary factor in assessing the cost of deliberation in the form of search: Deliberation will always cost less when done at a higher abstraction level. However, the disadvantage of higher abstraction levels is greater error in the planning, hence less benefit. In Chaffworld, the agent planned using a "map" that included only long chaff (longer than some threshold). See Figure 3. Experiments to test the effects of

abstraction are being designed. We also are designing experiments that will vary the abstraction level of the A* search by using finer and coarser step sizes in the representation.

Finally, we are addressing the question of the effect of uncertainty on the tradeoff between deliberating and reacting by designing an agent that is sensitive to the uncertainty of its environment and deliberates until the accumulated uncertainty makes further deliberation futile. We believe that such an agent will act effectively in both highly uncertain (high entropy) and less uncertain (low entropy) worlds. Furthermore, we believe that the agent will naturally act in a deliberative manner in low-entropy worlds and reactively in high-entropy worlds.

Essentially, we claim that the design of an agent must reflect the task it undertakes and the environment it experiences. In some environments, the agent will be reactive, while in others it will be deliberative. In environments that exhibit qualities of both reactive and deliberative environments, the agent design may well be a "two-loop" design. We can design the right kind of agent for its environment by knowing the cost/benefit tradeoff in the task, the level of abstraction of the representation used in deliberation, and the entropy of the environment.

3.3 Follow-up on Previous ONR-sponsored Research

3.3.1 A Textbook for Empirical Methods in Artificial Intelligence.

In the final report for our previous ONR contract we detailed several ongoing activities to broaden and disseminate our methodological approach: a Workshop on AI Methodology held in June 1991, the development of a curriculum in *agentology* during a modeling summer school for graduate students in 1991, and a AAAI-92 Tutorial on Experimental Methods for AI Research. These activities are culminating now in a textbook being prepared for use in graduate AI methods courses. Entitled "Empirical Methods in Artificial Intelligence," this textbook is a primer for the empirical evaluation of the new generation of agents being designed by AI researchers.

3.3.2 An Automated Model-Building Assistant.

Under our previous ONR contract (N00014-88-K-0009) we developed the Modeling, Analysis and Design (MAD) methodology (Cohen 1991), an approach to the design of autonomous agents that combines analytical modeling with empirical substantiation to justify the agent architectures we use. We are employing this approach in the context of the Phoenix system, a complex, real-time task environment in which autonomous agents fight simulated forest-fires (Cohen et al. 1989, Howe, Hart & Cohen 1990). With support from a recently completed URI award (N00014-86-K-0764) we developed a number of tools that support model-building and analysis, and have used these tools extensively in Phoenix (Cohen 1990a, Howe & Cohen 1991) and more recently in the domain of transportation planning. We are working to develop an *automated model-building assistant* capable of extracting causal relationships from data and constructing analytical models to guide the design process. The long-term goal of this research is to embed this system in an autonomous agent architecture,

creating agents capable of integrating perceptions of their task environments into conceptual models that support reasoning and planning. This work, outlined below, is described in more detail in (Silvey, Loiselle & Cohen, 1992).

The primary objective of the IGOR¹ project is to concentrate on understanding and automating the expert human process of incrementally building models from both experimental and non-experimental data. As such, this work represents the development of an enabling technology for principled AI research using the MAD methodology. IGOR is designed to function as a model-builder's intelligent assistant. The basic approach is to integrate the complementary strategies of exploratory and confirmatory data analysis (eg., (Tukey 1977, Wickens 1989)) in a knowledge-based decision aid. A comparable approach which integrates complementary data-driven and theory-driven strategies is described in (Langley et al. 1987).

In a broad context, model-building is the process of using observations of a functioning system to infer stable relationships between elements of that system. The inferred relationships facilitate prediction, explanation, diagnosis, and planning tasks; and they may be causal, qualitative, and/or quantitative. Models are built by agents, both artificial and human, to summarize, abstract, and generalize from experience. Since models represent knowledge, model-building is the process of learning or acquiring knowledge. Models are constructed and refined through a process of induction from observation; observation which may be active as well as passive. For example, the environment which contains or comprises the system in question is often occupied by the model-builder as well, and actions taken by that agent will themselves affect the problem environment. In many such cases, the resulting models provide a subsequent means for the agent to intelligently control the environment, as they consist of knowledge of the expected effects of agent actions in particular situations.

As humans, the models we construct from experience in our environment are based upon recognizable dependencies between a small number of observable variables. From this, we are able to construct causal explanations of how the system in which we are embedded, functions. The specificity, or level of abstraction, of a model can range from fully quantitative functional relationships, such as physical laws, to probabilistic or qualitative influence relationships, such as our intuitive understanding of economics. The abstract, higher level models are beneficial because they provide interaction constraints that hold for all more specific, lower level models. These constraints can reduce the search required to find specific, quantitative models, because much of the possible search space is ruled-out by the general patterns of dependencies in observed data.

The basis of constructing causal models from observed data is in our ability to detect conditional independency relationships. The methods we are using in IGOR to find causal structure in data are based on algorithms developed by Pearl (Pearl & Verma 1991, Pearl 1988). The approach is to first use tests for conditional independence to

¹ IGOR is so named because it is intended to be the MAD Scientist's assistant.

construct the adjacency structure of the model, an undirected graph in which the nodes represent observable variables and arcs represent direct dependency relations. This structure is then turned into a hybrid graph that describes an equivalence class of directed, acyclic graphs (DAGs) that correspond to possible causal models. This hybrid graph has arcs that have zero, one, or two arrowheads, roughly corresponding to undetermined, causal, and spurious correlations, respectively. The undetermined correlations are underconstrained in the data, their directionality could go either way, depending on the assignment of other unconstrained arcs. The spurious correlations represent hidden causes, or latent variables as they are sometimes called, that simultaneously affect the two observed variables.

These algorithms are based on the fact that the "true" causal model which governs the observed system imposes some constraints on the joint probability distribution over the observed variables. Therefore, we can use statistical tests on data samples of those observables to partially reverse-engineer the process. However, in general, it is not possible to construct a single causal model that is uniquely determined by the data. Additional, background knowledge may be required to take a partial model and complete it to the experimenter's satisfaction. This might include adding arrows to unconstrained arcs so as to agree with the flow of time, or the addition of specific latent variables that are known to the experimenter but were not measured in the specific data set being analyzed. Thus, we see these automated methods for discovering causal relationships as one tool, albeit a powerful one, for building models from data.

IGOR is being designed using the blackboard paradigm, where the shared blackboard data structure holds elements of the developing model, and the individual knowledge sources operate on those elements to create terms, hypothesize relationships, and perform statistical tests. The long-term objectives for IGOR include fully automated model-building and discovery mechanisms driven by an opportunistic control strategy. We expect to develop the automated strategies from our experience with the system as a manual analysis decision aid, letting human analysts provide the initial reasoning control strategy.

4 Publications, Presentations, Reports

Refereed Papers

Anderson, S.D., Hart, D.M. & Cohen, P.R., 1991. Two ways to act. *1991 AAAI Spring Symposium on Integrated Intelligent Architectures*. Also in *SIGART Bulletin*, 2(4): 20-24.

Silvey, P.E., Loisel, C.L. & Cohen, P.R. Intelligent data analysis. To appear in *AAAI92 Fall Symposium on Intelligent Scientific Computation*. Cambridge, MA. October, 1992.

Unrefereed Reports and Articles

Ravela, Srinivas S., 1992. A survey of reactivity. Technical Report #92-61, Computer Science Department, University of Massachusetts, Amherst.

Books or Parts Thereof

Cohen, P.R. *Empirical Methods in Artificial Intelligence*. Textbook in preparation (see Appendix A).

Invited Presentations

Paul R. Cohen

- Panel member, "The Future of Expert Systems" chaired by Dr. Y.T. Chen of NSF at the World Congress on Expert Systems, Orlando, Florida, December 16-19, 1991.
- Panel member, "The Empirical Evaluation of Planning Systems: Promises and Pitfalls" at the *First International Conference on AI Planning Systems* at the University of Maryland, June 16, 1992.
- Panel member, "Planning Under Uncertainty," *SIGMAN Workshop on Knowledge-Based Production, Planning, Scheduling and Control*, Workshop Program, *Tenth National Conference on Artificial Intelligence*, San Jose, CA, July 1992.
- "Methods for Studying Agents: General Issues and Specific Examples." Invited talk at the U.S. Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Washington, DC. April 6, 1992.
- "Three Examples of Statistical Modeling of an AI Program." Invited talk at the University of Texas, Austin.
- "Methods for Agentology: General Concerns, Specific Examples." Invited talks at Virginia Polytechnic Institute and the University of West Virginia.
- Invited presentation of EKSL research at the Institute for Defense Analysis's Information and Science Technology Advisory Group on Simulation in Washington, D.C., June 17-18 (Cohen is a new member of the Advisory group).

Contributed Presentations

Paul R. Cohen

- Welcoming address (untitled) at the *NSF/DARPA Workshop on Artificial Intelligence Methodology*. Northampton, MA. June 2, 1991.
- "Intelligent Data Analysis." Contributed presentation to be given at the *AAAI92 Fall Symposium on Intelligent Scientific Computation*. Cambridge, MA. October 23-25, 1992

5 Transitions and DoD Interactions

5.1 EKSL Research Presented at the Naval Research Laboratory

Paul R. Cohen was invited to speak at the U.S. Navy Center for Applied Research in Artificial Intelligence at the Naval Research Laboratory in April. His presentation, entitled "Methods for Studying Agents: General Issues and Specific Examples," covered research conducted under this and previous ONR contracts. The first part of the talk presented the methodological approach developed under our previous ONR contract. The second part described the application of these methods in our current research (including an example from Chaffworld, described above).

5.2 Transition of Methodological Tools to DARPA Planning Initiative

Since the end of the previous contract we have refined our methodological approach and begun to operationalize aspects of it for use as components of IGOR (see above). We recently proposed to develop and test several *experiment modules* for the DARPA Planning Initiative (PI)². It is difficult to design experiments to evaluate planners' performance, and even more difficult to model the factors that affect performance. Under this proposal we would develop modules to facilitate experiments to answer the following questions: How well does a component of a planning/scheduling system work, and how will changing the task, environment, or structure of the component affect its performance? We will begin by developing five modules, with more anticipated. The first five are for:

1. inducing causal models of the factors that affect the performance of plans and planners
2. detecting statistical dependencies between actions and selecting the most likely explanations of the dependencies
3. inducing Markov state-space models of the behavior of planners and simulated plans
4. inducing acausal functional relationships between factors, including those that affect the performance of plans and planners
5. assessing nonparametric confidence intervals around variables such as plan completion time and number of units delivered.

We must stress that although these modules all have mathematical foundations in statistics and probability theory, we are proposing full software packages, not merely a handful of statistical tests. Experiment modules will help the researcher in all phases of experiment design, data collection, exploratory data analysis and model-building. Furthermore, the focus of our modules is explanatory models of AI system performance, not simple statistical hypothesis testing.

² This contract, "Experimental Methods for Evaluating Planning Systems," will be funded to start in FY93. It will involve joint work with researchers at Colorado State University.

Once the theoretical underpinnings of these modules are worked out, we will test them in an extensive set of experiments with two systems, the Phoenix planner and TPSS, the Transportation Plan Steering System that we are developing for the PI. The experiments will be sufficiently extensive, thorough and documented to serve as case studies of the applications of our experiment modules, and will encourage other researchers in the PI to adopt the modules. In addition, the modules will be made available as part of the Planning Initiative's Common Prototyping Environment (CPE), fully integrated with the UMass CLASP³ statistical package that is part of CPE.

³ Developed under URI contract #N00014-86-K-0764.

References

- Anderson, S.D., Hart, D.M. & Cohen, P.R., 1991. Two ways to act. *1991 AAAI Spring Symposium on Integrated Intelligent Architectures*. Also in *SIGART Bulletin*, 2(4): 20-24.
- Arkin, R.C., 1988. Reactive/reflexive navigation for an autonomous vehicle. *Proceedings in Computers in Aerospace IV*. Conference of the American Institute of Aeronautics and Astronautics, Wakefield, MA. Pp. 298-306.
- Arkin, R.C., 1987. Motor schema based navigation for a mobile robot: An approach to programming by behavior. *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*. Pp. 264-271.
- Chrisman, L., 1991. Intelligent agent design issues : Internal agent state and incomplete perception. *Working Notes of the AAAI 1991 Spring Symposium on Sensory Aspects of Robotic Intelligence*.
- Cohen, P.R., Hart, D.M. & St. Amant, R., 1992. Early warnings of plan failure, false positives and envelopes: Experiments and a model. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Lawrence Erlbaum Associates. Pp. 773-778.
- Cohen, P.R., 1991. A survey of the *Eighth National Conference on Artificial Intelligence*: Pulling together or pulling apart? *AI Magazine*, 12(1): 16-41.
- Cohen, P.R., Hart, D.M. & deVadoss, J.K., 1991. Models and experiments to probe the factors that affect plan completion times for multiple fires in Phoenix. EKSL Memo #17, Dept. of Computer Science, Univ. of Massachusetts, Amherst.
- Cohen, P.R., 1990a. Designing and analyzing strategies for Phoenix from models. *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control*. Katia Sycara (Ed.). Morgan Kaufman. Pp. 9-21.
- Cohen, P.R. (Ed.), 1990b. *Working Papers of the Workshop on Artificial Intelligence Methodology*. Northampton, MA. June 2-4, 1990. Forthcoming.
- Cohen, P.R. & Howe, A.E., 1989. Toward AI research methodology: Three case studies in evaluation. *IEEE Transactions on Systems, Man and Cybernetics*. 19(3): 634-646.
- Cohen, P.R., DeLisio, J.L. & Hart, D.M., 1989. A declarative representation of control knowledge. *IEEE Transactions on Systems, Man and Cybernetics*, 19(3): 546-557.
- Cohen, P.R., 1989. Why knowledge systems research is in trouble and what we can do about it. Technical Report #89-81, Dept. of Computer Science, Univ. of Massachusetts, Amherst.
- Cohen, P.R., Greenberg, M.L., Hart, D.M., Howe, A.E., 1989. Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3): 34-48.

- Dean, T. & Boddy, M., 1988. An analysis of time-dependent planning. *Proceedings of the Seventh National Conference on Artificial Intelligence*. St. Paul, MN. Pp. 49-54.
- Gruppen, R.A., 1991. Planning grasp strategies for multifingered robot hands. *Proceedings of 1991 IEEE International Conference on Robotics and Automation*. Pp. 646-651.
- Gruppen, R.A. & Weiss, R.S., 1991. Force domain models for multifingered grasp control. *Proceedings of 1991 IEEE International Conference on Robotics and Automation*. Pp. 418-423.
- Hansen, E.A., 1990. A model for wind in Phoenix. EKSL Memo #12. Dept. of Computer Science, Univ. of Massachusetts, Amherst.
- Hart, D.M., Anderson, S.D. & Cohen, P.R., 1990. Envelopes as a vehicle for improving the efficiency of plan execution. *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control*. Katia Sycara (Ed.). Morgan Kaufman. Pp. 71-76.
- Howe, A.E. & Cohen, P.R., 1991. Failure recovery: A model and experiments. *Proceedings of the Ninth National Conference on Artificial Intelligence*. Morgan Kaufman. Pp. 801-808.
- Howe, A.E., Hart, D.M. & Cohen, P.R., 1990. Addressing real-time constraints in the design of autonomous agents. *The Journal of Real-Time Systems*, 1(1/2): 81-97.
- Howe, A.E. & Cohen, P.R., 1988. How evaluation guides AI research. *AI Magazine*, 9(4): 35-43.
- Kaelbling, L.P. & Rosenshein, S., 1989. Integrating planning and reactive control. *Proceedings of the NASA Conference on Space Telerobotics*. Jet Propulsion Laboratory, Pasadena, CA. Vol. II, Pp. 359-366.
- Khatib, O., 1985. Real-time obstacle avoidance for manipulators and mobile robots. *Proceedings of 1985 IEEE International Conference on Robotics and Automation*. Pp. 500-505.
- Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M., 1987. *Scientific Discovery: Computational Explorations of the Creative Processes*. The MIT Press.
- Payton, D.W., Rosenblatt, J.K & Keirsey, D.M, 1990. Plan guided reaction. *IEEE Transactions on Systems, Man and Cybernetics*, 20(6): 1370-1382.
- Payton, D.W., 1986. An architecture for reflexive autonomous vehicle control. *Proceedings of 1986 IEEE International Conference on Robotics and Automation*. Pp. 1838-1845.
- Pearl, J. & Verma, T.S. 1991. A theory of inferred causation. *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, J. Allen, R. Fikes, & E. Sandewall (Eds.). Morgan Kaufmann. Pp. 441-452.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman.

Powell, G.M. & Cohen, P.R., 1990. Operational planning and monitoring with envelopes. *Proceedings of the IEEE Fifth Annual AI Systems in Government Conference*, Washington, DC.

Ravela, S.S., 1992. A survey of reactivity. Technical Report #92-61. Dept. of Computer Science, Univ. of Massachusetts, Amherst.

Silvey, P.E., Loiselle, C.L. & Cohen, P.R. Intelligent data analysis. To appear in *AAAI92 Fall Symposium on Intelligent Scientific Computation*. Cambridge, MA. October, 1992.

Silvey, P.E., 1990. Phoenix baseline fire spread models. EKSL Memo #13. Dept. of Computer Science, Univ. of Massachusetts, Amherst.

Tukey, J. W. 1977. *Exploratory Data Analysis*. Reading, MA: Addison-Wesley.

Wickens, T.D. 1989. *Multiway Contingency Tables Analysis for the Social Sciences*. Lawrence Erlbaum Associates.

Appendix A: Textbook Prospectus

Empirical Methods in Artificial Intelligence

Paul R. Cohen

I. Introduction

- A. AI Programs as Objects of Empirical Studies**
- B. Versions of the Three Basic Research Questions**
- C. A Methodology for Achieving General Answers to the Basic Research Questions**
- D. Kinds of Empirical Studies**
- E. Data Analysis for the Four Classes of Studies**
- F. Empirical AI in Context**

II. Exploratory Data Analysis

A. Data

- A.1. Scales of Data**
- A.2. Transforming data**
- A.3. Measurement Theory**
- A.4. Causal Models of Data Values**

B. Visualizing and Summarizing Data

- B.1. Exploring One Variable**
- B.2. Statistics for One Variable**

C. Joint Distributions

- C.1. Joint Distributions of Categorical or Ordinal Variables.**
 - C.1.1. Dependencies Among Row and Column Variables in Contingency Tables**
- C.2. Contingency Tables for More than Two Variables**
- C.3. Statistics for Joint Distributions of Categorical Variables**
 - C.3.1. An Easy and Useful Special Case: 2 x 2 Table**
- C.4. Visualizing Joint Distributions of Two Continuous Variables**
 - C.4.1. Finding Hints of Causal Relationships in Scatterplots**
 - C.4.2. Point Coloring to Find Potential Causal Factors**
- C.5. Statistics for Joint Distributions of Two Continuous Variables**
 - C.5.1. Pearson's Correlation Coefficient**

D. Time Series

- D.1. Visualizing Time Series**
 - D.1.1. Smoothing**
- D.2. Statistics for Time Series**

- E. Series of Categorical Variables. Behavior Traces
 - E.1. Visualizing Behavior Traces
 - E.2. Statistics for Behavior Traces
- III. Fundamental Issues in Experiment Design
 - A. Some Terminology
 - B. The Concept of Control
 - B.1. Control Conditions in MYCIN: A Case Study
 - C. Four Spurious Effects
 - C.1. Ceiling and Floor effects
 - C.2. How to Detect Ceiling and Floor Effects
 - C.3. Regression Effects
 - C.4. Order Effects
 - D. Sampling Bias
 - D.1. A Sampling Bias Due to an Arbitrary Cut-off
 - E. The Dependent Variable
 - F. Pilot Experiments
 - G. Guidelines for Experiment Design
 - H. Summary
- IV. Hypothesis Testing and Estimation.
 - A. Statistical Inference
 - B. Introduction to Hypothesis Testing
 - C. A More Formal View of Statistical Hypothesis Testing
 - C.1. Sampling Distributions
 - C.2. How to Get Sampling Distributions
 - C.2a. The Sampling Distribution of the Proportion
 - C.2b. The Sampling Distribution of the Mean
 - C.2c. The Standard Error of the Mean and Sample Size
 - C.2d. Standard Errors for Other Statistics
 - D. Parametric Tests of Hypotheses About Means
 - D.1. The Anatomy of the Z Test
 - D.2. Critical Values
 - D.3. p Values
 - D.4. When the Population Standard Deviation is Unknown
 - D.5. When N is Small: The t Test
 - D.5a. One Sample t Test
 - D.5b. Two Sample t Test.
 - D.5c. The Paired Sample t Test

- E. Parameter Estimation and Confidence Intervals
 - E.1. Confidence intervals for μ when σ is known.
 - E.2. Confidence intervals for μ when σ is unknown.
 - E.3. An Application of Confidence Intervals: Error Bars
 - E.4. Hypothesis Testing with Confidence Intervals
- F. Some Practical Issues
 - F.1. Reporting Confidence Intervals, When and Why.
 - F.2. Errors
 - F.3. How Big Should N Be?
- G. Appendix 1
 - G.1. Degrees of Freedom
- V. Computer Intensive Hypothesis Testing
 - A. Monte carlo
 - B. Randomization
 - C. Bootstrap
 - D. Exact Nonparametric Tests
- VI. Testing One-Factor Models
 - A. Case Studies of Experiments
 - A.1. Addition Studies
 - A.2. Ablation Studies
 - A.3. Roughening Studies
 - A.4. Minimal Pair Studies
 - B. Data Analysis for One Factor Experiments
 - B.1. One-Way Analysis of Variance
 - B.2. Simple Regression
 - B.3. Tests on Learning Curves
 - B.4. Dependency Detection
- VII. Testing Multiple Factor Models.
 - A. Case Studies of Interactions between Architecture And Environment Factors on Behavior
 - B. Data Analysis for Multiple Factor Models
 - B.1. Two- and Three-Way Analysis of Variance
 - B.2. Multiple Regression
 - B.3. Log-Linear Analysis
- VIII. Building and testing causal models with path analysis.
- IX. Tactics for Generalizing Results
 - A. Finding Representative Problems
 - B. Replicating Results