DECIDING WHETHER TO PLAN TO REACT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Vlad Grigore Dabija

December 1993

I certify that I have read this dissertation and that in my opinion it is fully ac
in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

____
Barbara Hayes-Roth
(Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully ac
in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

____
Jean-Claude Latombe
(Co-Adviser)

I certify that I have read this dissertation and that in my opinion it is fully ac
in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

____
David M. Gaba

I certify that I have read this dissertation and that in my opinion it is fully ac
in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

____
Oussama Khatib

Approved for the University Committee on Graduate Studies:

____

Abstract

Intelligent agents that operate in real-world real-time environn have limited resources. An agent must take these limitations into ac when deciding which of two control modes - planning versus reaction - control its behavior in a given situation. The main goal of this thesis develop a framework that allows a resource-bounded agent to decic planning time which control mode to adopt for anticipated possible ru contingencies. Using our framework, the agent: (a) analyzes a com (conditional) plan for achieving a particular goal; (b) decides which c anticipated contingencies require and allow for preparation of re responses at planning time; and (c) enhances the plan with pre reactions for critical contingencies, while maintaining the size of the the planning and response times, and the use of all other critical resou the agent within task-specific limits. For a given contingency, the decisi plan or react is based on the characteristics of the contingency, the ass reactive response, and the situation itself. Contingencies that may occur same situation compete for reactive response preparation because c agent's limited resources. The thesis also proposes a knowledge represe formalism to facilitate the acquisition and maintenance of knowledge ir in this decision process. We also show how the proposed framework ca adapted for the problem of deciding, for a given contingency, wheth prepare a special branch in the conditional plan under development leave the contingency for opportunistic treatment at execution time. We a theoretical analysis of the properties of our framework and demonstrate them experimentally. We also show experimentally that framework can simulate several different styles of human reactive beh described in the literature and, therefore, can be useful as a basi describing and contrasting such behaviors. Finally we demonstrate th framework can be applied in a challenging real domain. That is: (a knowledge and data needed for the decision making within our fram exist and can be acquired from experts, and (b) the behavior of an age uses our framework improves according to response time, reliability resource utilization criteria.

Acknowledgements

I would like to thank my adviser Barbara Hayes-Roth for her signif support and guidance throughout my three years here at Stanford. T many long discussions, Barbara helped me clarify my thoughts. She always challenged my ideas in constructive ways and often prov alternative suggestions and views from different perspectives. M significantly, she always challenged me to look for more general solutic more meaningful results in my research. She taught me important lesso how to validate my work and how to present it. I have always appre Barbara's availability when I needed her advice most; on so many occ she found the time to see and advise me with very little or no notice at a but not least, she suggested lots of tips on how I can give my daughter a education.

I would like to thank my co-adviser Jean-Claude Latombe for his ad both while supervising my research and during the time I was his te assistant. His thought-provoking questions and perspective were very in enlarging the scope of my research. I am indebted to Jean-Claude f significant time pressure he exerted on me to complete my thesis during without which it could have taken me much more time to complete research.

I was very lucky to benefit from the advice and medical expertis Dave Gaba during my work here. I thank him for his patience and the t took to supply me with expert data and to give me his feedback on experimental results. In numerous occasions he commented extensively ideas from his perspective both as an expert in anesthesiology and an pilot. His computer science oriented comments were very helpful crystallizing many ideas of my work.

I would like to thank Oussama Khatib for his advice on my work. He pointed out related problems to my initial ideas, and challenged me to w solutions to them.

Serdar Uckun has provided me repeatedly with expert data in medical and car driving domains. I have also benefited from num discussions with him regarding my thesis or parts of it as well as other topics. I want to thank him for all his time and efforts.

I have had innumerable discussions with my office mate David Ash huge variety of subjects, related or not to this thesis. He was always a ahead of me in completing his Ph.D. program, and has always been a tou to follow. The fact that we had our oral examinations in the same da probably no coincidence. He has also provided me with a detailed desc of his reactive planner model, which I have used in some of my experim

I would like to thank to all the experts who have helped me with expertise in the domains in which I have conducted my experiments: Gaba, Serdar Uckun, David Ash, Alex Brousilovsky, Lee Brownston, Ja

Murdock, Rich Washington and Michael Wolverton. They have not
provided the domain data needed for the framework, but they hav
evaluated the system's performance.

My work has also benefited from numerous discussions with:
Brownston, John Drakopoulos, Mike Hollander, Philippe Lalanda, Jan
Larsson, Alex Macalalad, Philippe Morignot, Janet Murdock, Marcel Scho
Evangelos Simoudis, Henny Sipma, Rich Washington, Michael Wolverton.
have all contributed on many occasions their comments and advice.

I would like to thank Ed Feigenbaum, Rich Fikes and Bob Engelmore
supporting my work at the Knowledge Systems Laboratory. My thanks al
towards the administrative team of KSL: Grace Smith, Peche Turner
Michelle Perrie, for a wonderful job which has made any bureaucracy
things unknown to me.

Two other groups of people have given me wonderful opportu
during my years at Stanford, and I want to thank them here. Dr. Shogo
and Dr. Katsuhiko Tsujino gave me the opportunity to spend three wo
months in their research group at the Mitsubishi Central Research Lab
in Japan. Largely due to them, my work there was extremely productiv
my touristic experience was unforgettable. My visit there was the beg
of a wonderful friendship with them. Prof. G. Schweitzer and Na
Tschichold-Gürman gave me a similar opportunity in Switzerland, and I
like to thank them as well.

I would also like to thank Ranan Banerji and Val Breazu-Tannen
their support and encouragement, as well as their unlimited confidence
abilities. Their support made this entire endeavor possible.

Financial support for my work was provided by ARPA Grant NASA N
2-581: Intelligent Systems for Science and Engineering Applications.

Finally, I would like to thank my parents, Antoaneta and Constanti
their love, sacrifice and support in everything I have ever done, an
encouraging me to pursue my studies. The pride they derived from my
has given me the strength to complete this thesis. I wish to thank m
years old daughter Dominique (born between the theory and th
comprehensive exams) for being such a wonderful child and for the be
moments she gave me, which have recharged my batteries many times
my work on this thesis. Last but by no means least, I want to thank m
friend and wife Tatiana for her infinite patience during the long hours
into this work. She gave me the opportunity to begin this work and
supported me all along the way. She provided me with unlimited lov
support and the right balance of encouragement and restrain to ke
going through the most difficult moments. It is to these four wonderful
that I affectionately dedicate this dissertation.

To: Toni, Puiu, Tatiana and Dominique.

Table of Contents

List of Tables

List of Illustrations

Chapter 1
Introduction

How should an intelligent agent prepare to satisfy a goal, while b able to respond to the great variety of contingencies that might impe achievement of goals? Short answer: through planning. For a m comprehensive answer, you may want to read this thesis. It may provic with a partial answer to this question, but it may also raise many questions.

Many AI research resources have already been devoted to fin solutions to the problem of planning, usually defined as choosing the ne or steps for the execution of a system, based on knowledge of the p situation, the system's goals, and the operators available. The essen planning in AI is the ability to reason about actions and their effects equally important, this reasoning process can take place before the execution starts. Therefore, it must deal with all the uncertainties due fact that the actual situation at execution time can only be assum planning time, when many characteristics of the environment either be taken into account, or simply cannot be known. Many activitie Computer Science can be regarded as instances of planning. One exam programming, which requires making decisions (at planning - programming - time) about actions to be performed later, at pro execution time, based on expectations about the environment in whicl will be executed. A computer program is a formal specification of hov resources of the computer will be applied to solve a given problem. Alt conventional plans are not synonymous with programs, as also argu [Drummond, 1989], we briefly use the analogy here for explanatory pu The more complex and unpredictable the execution environment is, th contingencies can occur during a program execution. The programmer therefore prepare the computer to properly respond to as many of contingencies as possible, while still keeping the program within computer resources, that is, it must still be small enough to fit in memo must still be fast enough to give an answer in a required amount of tim same situation occurs in all other domains in which planning is require

A special kind of planning is reactive planning, i.e. building a se specific perception-action rules stored in a computationally efficient [Brooks, 1986; Agre & Chapman, 1987]. From now on, we will call this ty planningreaction, as opposed to the conventional type of planning whic will call simply planning or sometimes, to clearly distinguish it from react conditional planning. To continue our parallel with computer progra interruptions, traps, exceptions, and error treatment routines in a p can be regarded as reactions. They are executed as response to a large of specific situations, and are not necessarily intended to ensure the su normal continuation of the program towards completing its final Sometimes, they are just intended to allow the program to interact gra with the environment or to help the program recover from a critical and allow the user to intervene to facilitate the continuation of the p or maybe to start the execution of another program, or even to write program (to replan).

All the characteristics discussed so far for computer programm apply to most domains where planning is needed as a means of en proper behavior of the system, before starting the actual execution c system to achieve a given goal. Such domains range from "high-le cognitive, symbolic domains like medical fields (e.g. anesthesiology, int care monitoring), to "low-level" manipulation domains like robot mani control. The common characteristics of all such domains is that their p tasks can be (at least conceptually) translated into computer program therefore conform to our previous discussion.

The same planning problem can be of very different levels of diffi depending on the assumptions made about the environment in which t is to be executed. For a well structured, "well behaved" environment will not present "surprises" to the executing agent, the planning prob much easier than for a more natural environment. In the latter case, contingencies are possible during plan execution. We will call a contin any state of the world entered by the executing agent while following a that should not have occurred as a result of executing the plan up to tha Contingencies are the effect of interactions between the agent anc environment; they occur because of: (i) predictable actions of environment, or (ii) the unpredictability of the environment, or (ii unpredictability of the execution subsystems of the agent. In the real the number and variety of contingencies that can occur during the ex of a plan is unlimited. An ideal planner should take care of all contingencies and build a "universal" plan [Schoppers, 1987] for the age has already been shown [Ginsberg, 1989], building such a plan is not f for interesting application domains, due to practical limitations of the resources. However, many of these contingencies can be ignored, e because they do not seriously affect the execution of the plan or becau have an extremely low likelihood of occurrence. Some of the rema contingencies may have a very high likelihood of occurrence while requiring elaborate subplans to treat them. Therefore, these subplans be included as conditional branches in the original plan. Other signif less likely contingencies may allow for a very short time of response, having disastrous consequences if the response does not occur in time contingencies probably should be treated reactively. These reactions ne lead the agent to the final goal of the initial plan; it is enough if the stabilize the situation, avoid the consequences of the contingency, and the planner to replan a comprehensive solution from the current situa the final goal. Yet other contingencies, not extremely likely and without term dramatic consequences, can be ignored at planning time and left possible replanning phase at execution time: when they appear, the (which is not under very high time pressure) can suspend execution ar its time to replan a solution from that situation to the final goal. Thi involve either a complete solution or, more frequently, a patch to bri agent back to one of the states in its original plan from which it can co execution (one such mechanism was implemented by the triangle table in STRIPS [Nilsson, 1984]).

From the above discussion we can derive the two basic control mod an agent that must deal with such contingencies: planning and reacti planning we will understand here both building a course of action before starting its execution and dynamic replanning, i.e. interleaving pl with execution. Each of these two modes has its advantages in ce circumstances, and we shall summarize them here. [Hayes-Roth, 1993] p a complete discussion of these characteristics.

Among the strengths of the planning model is the fact that plans c built to have a set of desirable global properties regarding the goals attained and the resources of the agent. The side effects of the actions executed as part of the plan can be carefully taken into account and a before execution begins. These properties are achieved by taking into complete descriptions of the states of the world as they are predicted planner. Of course, these states will conform to reality only if environment behaves according to the model that the planner has ab The more incomplete this model is, the more uncertainty in the beha the environment, and the more uncertainty about the actual states that encountered by the agent during plan execution. The final plan has degree of coherence and is easily comprehensible by a human user (th point is very important in domains where the entire credibility of the depends on how much a user can understand from the reasoning system, such as medical domains). The plan generated by a conditional usually makes a close approximation of the optimal usage of the a resources. Finally, the planned actions can be executed promptly at ru (since the agent simply follows a completely specified plan, in which th action is taken according to the plan, maybe after evaluating the res some tests in the case of conditional plans). However, the planning mo its weaknesses with respect to the real world. The two main disadvantag (i) the high computational cost of planning (which makes it necessa carefully consider which contingencies should be exhaustively treated i way - otherwise the time to build the plan may become prohibitive); a the inflexibility of the planned behavior - the agent can only act in st the world which are specified in the plan, and its performance will d very abruptly with any variations to such states.

The reactive model constructs a set of goal-specific perception-a rules and stores them in a computationally efficient form. The advantages of the reactive model are its flexibility of response to a larger run-time conditions (since each response is less carefully analyzed th the previous case, and the response does not need to embody a co solution to the final goal but can merely be an action to stabilize the s and allow the time for replanning) and its short time of response (det by the efficient way of storing the reactive plan). On the other hand, r still cannot anticipate, distinguish and store all runtime contingencies. therefore still exhibit precipitous failure in unanticipated conditions. B main disadvantage of reaction is that it is taken after a superficial eva of the current state, and does not benefit from an in depth analysis of t and the related action consequences. Therefore, while a reaction m locally appropriate, its global effectiveness is uncertain.

The planning and reactive control modes are near the end-points theoretical continuum of control modes. Together with two other c modes (reflex and dead-reckoning), they form a two-dimensional described in [Hayes-Roth, 1993]. Also analyzed there is the correspon between the space of control modes and a two-dimensional space of c situations, as well as the effect of combining the control modes in di degrees on the quality of run-time behaviors in the corresponding sp control situations.

We believe that planning and reacting complement each other, therefore we envision agents that: (a) plan courses of action design achieve goals under certain anticipated contingencies - conditional br are built in the plan for the very likely contingencies that also re significant planning to reach the goal; (b) augment these plans context-dependent reactions for noticing and responding to less likel important exogenous events; (c) control their behavior by following plans, while simultaneously monitoring for and, when appropriate, ex reactions associated with particular phases of their plans; and (d) revis plans when local reactions do not adequately address unanticipated eve

However, this complementarity of the planning and reaction co modes in intelligent agents is overlooked by many researchers today. planning research to date has been concentrated either towards just one two control modes, or when it attempts to combine them, the main purpo increase the reactive capabilities of the agent and to unload the conve planner's responsibilities. In this latter case, the general assumption i reaction comes for free, that is, either the agent's resources are unlim the reaction process does not use any significant amount of the a resources. Unfortunately, this is not the case in reality: any real age limited resources, and the reaction process may use significant amounts agent's resources. This fact has a couple of consequences: (i) a decrease reactive responsiveness of the agent (or equivalently an increase i response time to a given contingency), which may make some reac useless if they come too late, and (ii) a limitation in the number of re for which the agent can prepare in a given situation. This means th agent must be more selective in the types of reaction it prepares for situation, preparing the most important reactions and ignoring the oth the following chapters we define and characterize the value of reactio identify those characteristics of the agent and its working environme influence the response capabilities of the agent to different situations may encounter in its working environment. Based on this analysis formulate a framework to decide, at planning time, which control mo choose for contingencies that may appear during plan execution, that framework to decide, at planning time, whether a certain situation r special preparation for a possible reactive response, or whether it can for dynamic replanning at execution time. The problem is partic important for planning the activity of an intelligent agent which mus in a dynamic, complex, unpredictable real-time environment.

The approach begins with a plan designed to achieve a goal enhances it to cope reactively with critical contingencies, while maint the size of the plan and the planning and response times within rea limits. The framework can also be modified for the problem of deciding, given contingency, whether to prepare a special branch in the (condi plan or to leave the contingency for opportunistic treatment at executic

As an example, consider driving a car between two given locati Before starting, the driving agent plans its route in some detail, inc turns at intersections and expectations of achieving milestones along th in order to minimize travel time. It also prepares a conditional branch plan as an alternative route in case the original route is blocked at a intersection where blockage is highly probable. This conditional b requires extensive planning resources but produces a complete solutio leads all the way to the final goal. Along the way, the agent in fact enco unexpected heavy traffic and revises the remainder of its plan to ta alternate route. As it follows the revised plan, the agent passes a school, it watches carefully for children who might suddenly run into the street. As leaves the neighborhood of the school and enters an industrial area, th forgets about children and watches for other contingencies (e.g., railway crossings, trucks coming out of driveways). Note that the agent, executing the plan, is prepared to react to certain contingencies at d stages of the plan, while using dynamic replanning to solve o contingencies.

Given certain conditions (like the time of day, the weather, the ty roads to be used) the agent prepares in advance for possible conting that may appear on certain portions of the trip. However, it does not expectations for and responses to these contingencies as steps of the since they are not essential for the plan to be executed successfully. O other hand, if they happen and are not responded to properly, the preclude the successful completion of the plan. Examples of s contingencies are: sliding on a slippery road in cold weather, an unsi object in the street during night time, a child running in front of the c a nearby school, or a traffic jam at rush hours. Note that these contin were qualified by the characteristics of the situation in which they are to appear. For some such contingencies, a reactive response must alread since the situation does not allow enough time for the agent to rep solution. There exists an infinite set of such contingencies, so the cannot prepare to always react to all of them. Moreover, due to li computational and non-computational resources, if the agent prepares large a set of contingencies in a situation, selecting the correct respon the one that actually occurs may become a too long process, thus ren the response ineffective. However, the responses to such contingencies need to include an entire solution to the main plan's ultimate goal; if th responds to them fast enough to avoid unwanted consequences, then take the time to replan the entire solution from there on. Since contingencies are too many and not very likely, they do not warr complete conditional branch in the initial plan to lead to the final goal.

Therefore, we need a decision framework to guide the selection contingencies for which a reactive response should be prepared at pl time. This need arises in many domains besides car driving (for examp intensive care monitoring, anesthesiology [DeAnda & Gaba, 1991; Fish 1991; Gaba & al, 1991; Gaba 1991], nuclear power plant operation [Woods 1987]). Formulating this framework is an important step toward buildi control engine of real-time intelligent agents with limited resources fo domains. The formulation and evaluation (theoretical and experiment such a framework is the topic of this research.

In the following chapter, we outline the problem in more precise t We define the notion of contingency and classify contingencies into according to their importance and the way they should be treated by th (with conditional plans, with reactions, or simply ignored at planning and left for dynamic replanning if necessary). We also characterize domains where the framework developed here will be most applicable. a review of related work points out similarities with other paradigms.

Chapter 3 presents the basic approach. After giving an intui solution for a simple problem in the driving domain and analyzing solution, we present the details of the framework for the reaction prep decision. We show how it can be used to establish the value of reacting contingency in a given situation and to make the decision of whether t to react to that contingency. The chapter closes with a discussion of ho framework may be modified and applied to decide whether a ce contingency, in a given situation, requires preparation of a complete in the initial conditional plan.

Chapter 4 discusses a proposal for a knowledge representa formalism for contingencies, reactions and situations, to facilitate structuring of the planner's knowledge and its manipulation.

Chapter 5 presents a theoretical analysis of the framework present chapter 3 for deciding whether to plan to react to a given contingenc given situation. A few formal properties are stated and justified, to s claims of generality and optimality (in terms of using the agent's reso for the proposed formalism.

Experimental demonstrations are then presented and briefly ana in chapter 6. Three domains were used for this purpose: an everyday where everyone is an "expert" (car driving) and two highly specia medical domains of expertise (anesthesiology and intensive care monit Results include simulations of several models of human reactive beh discussed in the literature. A demonstration in a complex, real-w application domain shows: (1) that the knowledge and data needed f decision making process exists and can be acquired from experts in domain; and (2) that the behavior of the agent improves (accordi response time, reliability and resource use criteria) as a result incorporating our decision framework in the agent's planning mechan

After summarizing our work, we make in chapter 7 a few suggestio natural continuations of this research, including applications of case reasoning techniques for managing a library of reactive plans and a of contingencies and reactions, and several applications of lear mechanisms to different parts of our framework.

Appendix 1 briefly presents the architecture of the reaction de module and the interface for integrating the module in an intelligent a

Appendix 2 completes the vocabulary example started in chapter the driving domain. It presents a large enough grammar to represent the situations, contingencies and reactions used as examples from this throughout the thesis.

In appendix 3 we present the results of a number of experiment have conducted in the anesthesiology domain, in order to provide evidence regarding the generality and applicability of our framewo real-world domains.

Finally, appendix 4 complements the presentation of intensive monitoring domain experiments in chapter 6, by presenting a few co sets of contingencies as they were ranked by our framework.

Chapter 2
The Problem

In this chapter, we outline the problem in more precise terms. define the notion of contingency and classify contingencies into three according to their importance and the way they should be treated by th (with conditional plans, with reactions, or simply ignored at planning and left for dynamic replanning if necessary). We also give a character of the domains where the framework developed here will be best app and what its limitations are. Finally a review of related work points similarities with other planning paradigms.

2.1. Contingencies

Let us consider first a more detailed version of the example present the previous chapter. Suppose the agent commutes each morning by ca home (starting point S) to the office (final goal G), as shown in figure 2 will limit ourselves to the study of a small segment of the car's route b points A and E. Suppose the route comes to an intersection with a traffi (point B). The fastest route between B and E is through C, which is the the agent normally takes if the traffic light at point B is green. Howeve driving agent knows that, if this traffic light is red, then many other lights between B and E through C will be red when the car will reach thus making the journey very slow. In the same time, the agent knows t at point B, it will take a right turn and go through point D, then it can point E (and therefore the goal G) much faster.

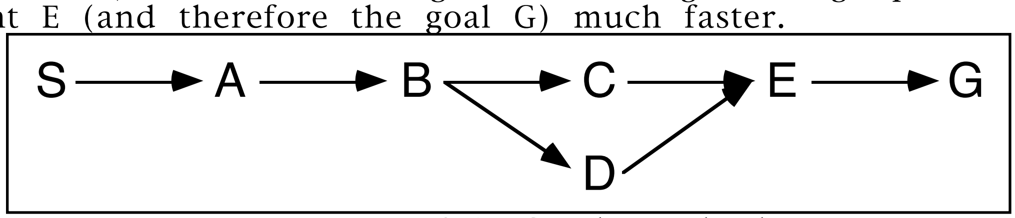

Figure 2.1. Conditional plan

The fact (and its associated state of the world) that the traffic light when the agent reaches the intersection at point B is a contingency, sinc not a result of the execution of the plan. In this case, the agent prep complete branch in the conditional plan to treat this contingency.

Suppose now that the point A in the plan built by the agent is a s in front of which the agent passes with its car. If the commute takes plac time when children are at school, or go to school, the agent prepares to carefully for children who might suddenly run into the street. It also that in front of a school, a ball may suddenly pop up in front of the car and many other contingencies (some more of which will be considered demonstrations described later on) may appear during the time when is in the school zone A. As it leaves the neighborhood of the school and another area (e.g. an industrial area), the agent forgets about childr balls and prepares for other contingencies (e.g. railway crossings, t coming out of driveways, etc.).

Let us consider for a moment the following three contingencies w appeared in the previous example: the traffic light at point B, the running into the street in front of the car, and the ball popping up in the car.[1] The common characteristic of these three contingencies is that are not generated as a result of the execution of the plan. We defi contingency to be any state of the world entered by the executing agen following a plan, which is not: (i) a direct consequence of executing actions of the plan up to that point, or (ii) an exogenously generated s the world assumed in the design of the plan. Therefore, a contingency d necessarily affect the agent or the plan execution, and when a contin does affect the plan, it is not necessary that it will negatively affect i example, a contingency may be a state which is not the current expecte according to the plan execution, but is a state which should have been along the way, after executing some more steps of the plan. The agen detect it and use it to skip the unnecessary steps in the plan, for exampl same way as it was done with triangle tables in [Nilsson, 1984]. To simpli exposition, from here on we will use the term contingency to also mean any fact or sign that was not expected as a result of the plan execution, and may indicate that a state is a contingency according to the previous def

The three contingencies presented above are very different in n and will be treated differently by our agent. The traffic light contingen happen very often (the actual probability to encounter a red signal is by the length of time the signal is green divided by the length of time i the signal to complete an entire cycle, provided that the signal is correlated with another signal previously encountered by the car and t signal behaves independently of the amount of traffic that passes thro for a two-way signal equally divided between the two directions of traffi probability is almost 0.5, though somewhat less because of the color yell likelihood of occurrence is significantly (one or more orders of magn higher than that of the other two contingencies. The treatment of contingency (by following an alternate route through point D to reach p and then the goal G) also needs an elaborate plan which must be prep advance (otherwise, after turning right at the traffic light, the agent stop and replan its route by possibly using maps, which may take a enough time to wipe out any savings obtained by avoiding the traffic lig the path through C). Therefore, the agent must prepare a conditional in the main plan for this contingency. This will use significant pla resources, but will have all the advantages associated with the pla control model discussed in the previous chapter.

The contingency defined by the child running in front of the c much less likely to happen than encountering a red traffic light, even driving in front of the school. This contingency has also a much h uncertainty about when and where it can occur. Thirdly, the plan to tr contingency is much simpler (it is usually enough to brake and maybe t to the right, depending on the distance to the child); after taking corrective action and avoiding the collision, the situation does not pres more dangers, so the agent can take its time to replan a course of acti will get it from the new state to the goal (this may be as simple as res the car, or as elaborate as finding an alternative means of transporta the car was damaged by hitting a pole on the side of the road while a the child). While the critical situation was avoided by a simple plan, th obtained after its execution is unknown and may belong to a large set o different states. Therefore, a comprehensive conditional plan to exhau treat all these states and preplan the agent's execution from them to th goal G may be prohibitive. The practical alternative is to treat contingencies in a reactive manner, by attaching simple reactive pla those points in the main plan where such contingencies may occur. Aft reaction will yield a non-dangerous state for the agent, it can take its t dynamically replan for a complete solution.

The third contingency stated before - the ball popping up in fro the car when driving along a school - is a little more probable than th running in front of the car, but the likelihoods of the two contingenci roughly of the same order of magnitude. However, in this third case consequences of hitting a ball with a car (especially with a relatively moving car in the vicinity of a known school) are significantly smaller in the child case. Moreover, the side effects of making a dangerous ma to avoid the ball may outweigh by far the consequences of hitting the Therefore, for such a contingency, the agent is much better off if it ign at planning time, thus conserving its limited resources for other important contingencies.

To summarize the discussion in this section, we have identified types of contingencies that may appear during the execution of a plan are classified according to the action taken by the agent at planning t prepare for their occurrence at execution time. These types of contin are:

- (i) contingencies for which the planner builds complete conditi branches, from the contingency state to the goal state, in the main
- (ii) contingencies for which the agent prepares reactive responses; are combined into reactive plans by a reactive planner, and attached to appropriate segments of the complete plan provided b conditional planner;
- (iii) contingencies ignored by the agent at planning time, either be their treatments can be left for dynamic replanning when they encountered at execution time, or because they are considered important than the contingencies included in the previous categories, and the agent simply does not have the resources to p a reaction (much less a complete branch in the plan) for them.

The justification for this classification is mainly related to the li resources that a real agent can use. For a few contingencies, the agen generate complete plans and combine them in a conditional plan. Ho the agent's limited planning and execution resources do not allow for to contingencies to be treated this way. Still, the agent can prepare at pl time reactive responses for a larger set of contingencies; these response not ensure full solutions to the goal state, but they will give the age possibility to dynamically replan its actions at execution time. But in n can a real agent with limited resources prepare for all possible contin in a real world application domain. Many of these contingencies mu ignored at planning time. The problem addressed in this thesis is how to which contingencies to select for preparation of reactive responses, which to ignore at planning time.

2.2. Summary of the Problem

The example problem outlined in the previous section highlights aspects of the general problem with which we are concerned. We shall here this problem more precisely, and then we will propose a solution fc the next chapter.

In all our previous discussion we have referred to reaction as planni a conscious form of preparing condition-action behavior. That is, the consciously prepares, before starting the actual execution, a set perception-action rules for a certain segment of the plan. They are executed by high level execution mechanisms of the agent similar to tho execute the main plan, and are not intended for execution by a "lower higher priority execution mechanism which may be part of the architecture (like the one proposed by [Brooks, 1986; Kaelbling, 19 Actually, the agent will resort to a reaction to a contingency only if it l conditional branch in the plan at that stage during the execution, an consciously take the decision to try to use reaction in that situation. Th not mean that we specifically prohibit in our agent architectures any level execution mechanisms which have the ability to react faster anc higher priority to certain contingencies. It only means that we are concerned with such precognitive types of reaction (e.g. locomotion reactions like avoiding obstacles by a moving robot). We are only conc here with contingencies to which such reaction mechanisms cannot re On the other hand, if the agent architecture does not include such lov reaction mechanisms, then the contingencies to be treated by them ma the set of contingencies which are analyzed by the higher level cog mechanisms of the agent using the framework proposed in this work.

Since we will talk more in the following section about th characteristics of the domains in which this work is best applicable, We simply say here that we are particularly interested in planning the acti an intelligent agent with limited resources and multiple goals working dynamic, unpredictable, real-time environment. The agent must itself real-time, i.e. be "predictably fast enough for use by the process serviced" [Marsh & Greenwood, 1986]. In order to behave properly, the must plan its actions ahead of time, and then monitor the plan executi be prepared to respond to contingencies that may appear during execution. This emphasizes two orthogonal qualities that the agent exhibit: sensitivity to run-time contingencies and commitment to sp goal-oriented actions. Such behavior can be accomplished by combinin two fundamental control modes mentioned before: planning and reactin

As will be shown in section 2.4, most research to date is conce either with employing only one of these control modes, or simply atten turn a system to become increasingly reactive and rely as little as possi planning. These works concentrate mainly on how to prepare rea responses and tend to use them in such a way as to substitute re planning. Our approach differs from these others in its recognition c complementary strengths and weaknesses of the two modes, and in it integration of planning and reacting within a single agent.

Our premise is that, whenever time and other resources allow dynamically planned response is never worse (and usually better responding to a contingency than a reactive response[2] previously prepared for it. There are several reasons for this assumption: (i) the repla response is generated at execution time when more information is availa opposed to planning time, when the reaction is prepared; (ii) replanning, an agent has time to analyze all the relevant information search for the best available solution by planning a complete solution the goal, while in order to react, the agent may have only a few alter (in the reactive plan) to choose from and only a few tests to decide response, which must therefore be taken based on incomplete infor obtained from an incomplete analysis of the current situation; (iii) if tim limited that it cannot even perform all these tests, the agent may have to more general action hoping to improve the situation at least temporar to buy more time to look for a better solution. The reason we need t reaction is that the replanned solution may be found too late and there of no more use at the time it can be taken. Thus, we assume that the im of regular planning makes it irreplaceable (due to the vast diversi situations in real-world environments), but the agent's real-time perfo can be significantly improved by preparing reactive responses for a l number of critical contingencies that may be foreseen to appear execution of the plan already built to achieve the main goal.

By not including enough contingencies for reactive treatment, performance of the agent will be suboptimal. On the other hand, by in too many such contingencies, the reactive response time becomes too thus degrading the system performance once again.

Unless otherwise stated, we assume that, given a contingency, the knows of an action (maybe a small sequence of elementary actions) wh applied reactively, either solves the problem generated by the continge at least postpones its deadline long enough to allow for replanning entire solution.

The main issue for us then is to enable the agent, for each phase main plan, to select the right set of contingencies for which to pr reactions. That is, our problem is to specify a decision framework which:

m given:
- l an intelligent agent with:
  - G capabilities:
    - F planning and dynamically replanning
    - F monitoring
    - F reactive behavior
  - G constraints:
    - F limited resources
    - F real-time performance
- l a (possibly conditional) plan by which the agent can achieve current goal
- l a set of contingencies known to possibly appear at certain ti during the plan execution, each with:
  - G reactive responses associated with them
  - G known characteristics associated with each such contingency gravity of consequences, time deadlines) and with their reac (e.g. resource requirements)

- enable the agent to decide at planning time on how to select a "rational" subset of these contingencies (according to a desired behavior pa for which the reactive responses should be attached to the main (while preserving the real-time responsiveness of the agent to all contingencies, given its limited resources).

We have used the word "rational" in the previous definition, ar needs some disambiguation. A behavior of the agent in a given situat defined by the order in which the agent classifies the set of contingenc that situation, according to the value of reacting to them. For the situation and set of contingencies, there are different behaviors that th may exhibit. Some of these behaviors may either not be suitable for situation, or may even be considered abnormal, hazardous or pathological. But there is at least one such behavior which is consi appropriate or normal for that situation, by the experts in the domair even possible that there are several different behaviors that ma considered appropriate in a given situation. Each behavior is appro according to a behavior model, and in the literature there have been de number of such reactive behavior types for domains in which critica stressful situations are common and very dangerous like aircraft flying 1991], nuclear power plant management [Woods & al., 1987] or anesthesia & al., 1991]. In most of the thesis we will refer to what is considered to "normal" behavior by experts in each domain from which we draw examples. However, in section 6.3, we will discuss some other types behaviors and how they can be translated and simulated with our fram

One problem related to the one we stated before is conditional pla As discussed before, there are three courses of action that an agent can prepare a response to a possible contingency: plan a conditional branc a reactive behavior, or ignore the contingency at planning time. Our a will focus on how to decide whether to prepare a reactive response contingency, but the general framework which will be developed for purpose is also applicable (with certain modifications) to the proble deciding whether to prepare an entire conditional branch in the mai for a possible contingency. In section 3.5 we will briefly discuss what a changes that must be made to our formalism so that it can also be used t which is the set of contingencies for which conditional branches shou planned. However, in the rest of the thesis, we will assume that the age already built the complete conditional plan, and is only trying to augr with reactive responses to as many contingencies as possible being limi its finite resources.

The selection criteria which we are looking for are much more cor than any utility measures (e.g., [Minton, 1990]) proposed so far. For ex in our approach, some of the contingencies associated with a situatio appear in practice with a very low probability, but they may be very cri they occur, and thus are worth preparing for reactively and are also being remembered. This is in contrast with most of the research to date, is mainly concerned with improving the systems' performance by ca into reactive plans the responses to the most frequently occu contingencies.

But before reviewing the previous research in this domain, let attempt to characterize first the domains in which the problem stated significant and where our solution framework is applicable.

2.3. Application Domains

Much of the planning work to date has concentrated on applicatio artificial domains. Such domains are well-structured and well-defined t system designer, which usually means that the entire set of pos contingencies is known in advance, and that this set is of a manageabl The main implication of this is that the resource limitations of the age be ignored (particularly at execution time) with respect to the size of th whether the main control mode employed is conditional planning or r planning, that is, we can always assume that we have a powerful enough to be able to respond in time to any of the contingencies that it knows This is clearly an artificial assumption which drastically simplifies planning problem and limits the applicability of the solutions proposed.

By contrast, we are interested here in applying the planning para to real-world domains and to allow the agent to operate in real-world closed and limited for practical purposes) domains. The main characteri such a domain and the agents operating in them is real-time defin [Marsh and Greenwood, 1986] as “predictably fast enough for use b process being serviced”. This means that the agent must be guarante respond, at execution time, in a prespecified time limit to any continge which it has prepared a response at planning time. However, if an ager limited resources prepares to respond to too many contingencies in a situation, than it may not be able to guarantee a timely response to th time-pressured of these contingencies: e.g. it make take too long for the to discriminate among the possible contingencies for which it is prepa react, from the time it detects a contingency and until it has to tak corrective action. An example of an interesting domain for our framew the car driving domain, which will be used for exemplification throu most of the thesis. If a child appears in front of the car at small distanc is very little time for the agent to discriminate among the contingenci which it is prepared to react in that situation and to decide what k contingency this is and how to react to it. For an agent with lin computational resources it may be therefore better not to prepare to r the same situation for a much less critical contingency like a ball com front of the car, or a sudden loss in the radio signal, and so on.

These observations are valid in real-life domains because anothe their characteristics: they are very large, both in the number and var contingencies that may appear (which has been noticed a long time a [McCarthy, 1977] when describing the qualification problem), and in variety of corrective actions that may apply. Each corrective action app to a certain contingency may be better suited in some situation th another one. Therefore, we will always consider pairs contingency-sit associated with each situation in which that contingency may arise a which that response is the best to this contingency. For well-struc (usually artificial) or very limited domains where the number contingencies and responses is limited, the framework described in this is not necessary, since it is conceptually possible to use a more powerful which can take care of all the contingencies in each situation.

As seen before, real-world environments are usually unpredictable, that is contingencies may occur at any time, or at least uncertain in that the effects of actions and the actual state of the world after the execution of step cannot be foreseen with utmost precision. Such domains are also dynamic in the sense that the state of the world may change withou participation of our agent, for example, as a result of actions of ot cooperative or antagonistic - independent agents working in the environment (e.g. there are other agents driving cars on the same str our agent and their paths may intersect[3]). In real domains some contin tend to appear associated with certain plan steps or situations an likelihood of their appearance may be different for different situations, others can appear at any time with the same likelihood. For example, always possible for a child to run into the street, or for a meteor to fall i street or for the car to fall to pieces, but it is impractical for the agent to the lookout for all of these possible events all the time. Real-world do also present a huge variety of situations. In each situation different contingencies can happen, and the same contingency may be vi differently in different situations. In certain situations, some conting are more likely or more important than others. If the agent has to dr car on a mountain road in winter, it should expect bumps or damaged of the road, or slippery roads, instead of, say, traffic lights. The agent prepare for yet another set of possible contingencies in the case of driv freeways. Also, the most effective responses associated with a conting which may appear in different situations may be situation dependen agent should therefore be able to selectively prepare itself for the critical contingencies in each possible situation along a prepared plan.

We should also note that some of the contingencies associated wi situation may appear with a very low probability, but they may be critical if they occur, and thus are worth preparing for. This is in co with most of the literature to date, since most authors are mainly con with improving their systems' performance by caching the most freq used plans.

We also assume that short plans (a single action or a small sequen actions), if applied reactively, are usually enough to either solve the p generated by the contingency, or at least to postpone its deadline long to give the planner the time needed to dynamically replan the entire under the new circumstances.

Most real domains which have the features described above are u characterized as high level, knowledge intensive domains. Examples of domains are some medical domains (e.g. intensive care monito anesthesia), nuclear power plant operation, aircraft flying, car driving on. These are contingency-intensive domains, in which many conting can appear and in which some of these contingencies are very time-c and / or with very high consequences, even if they do not appear wit high frequency. Although these domains also require (some more others) significant skill development (by skill we mean here auton low-level, unconscious reflexes to certain contingencies), their r characteristic is that the process of planning and responding to contin is knowledge-intensive and thus uses significant high-level cogni resources of the agent. Our framework can be in principle applied to domain, but its value and effectiveness can be questioned for very structured, artificial domains (like the blocks world) and for low-level, intensive domains (or such tasks in higher-level domains), like locon tasks (e.g. reflex obstacle avoidance) or fine-motion robot manipulation (e.g. the peg-in-the-hole insertion problem), in which the number diversity of contingencies is limited and well-known in advance.

Even for such limited but real domains, we can argue that framework can be applicable as long as the resources of the agent in are not powerful enough to completely remove the uncertainty in the An example of such a domain is robot motion planning. The main pr here is the uncertainty, at execution time, in the position and orienta the parts and of the robot (e.g., a manipulator) in the workspace. A cl planning methods developed for this problem deal with such uncertaint second phase of planning; in the first phase, plan skeletons and strategies are produced, using path planning methods which assume uncertainty (i.e. no contingencies) [Latombe & al., 1991]. Then diff methods are used to deal with contingencies generated by the aforeme uncertainties. For example, SPAR [Hutchinson and Kak, 1990] adds verifi and local recovery plans to reduce uncertainty and to prepare for p failures. Similarly to the reactions used in our framework, these recovery plans are only single, special-purpose actions (which may entered by the user) and are associated with uncertainty-reduction g priori. An inductive learning technique is used by [Dufay and Latombe, a trainer module generates patches to be inserted in the ground plan are local strategies refining the ground plan, similar to our reactive attached to the main plan (e.g. rotate a card to insert it into a slot). The further provides for the graceful degradation of its performance by a for entering rules on line if everything else fails. However, the most co technique for dealing with uncertainty-generated contingencies in domain is skeleton refining [Lozano-Perez, 1976; Taylor, 1976]. A skeleton (or assembly description) appropriate to the task at hand is retrieved a plan and then iteratively modified by inserting complements (e.g. readings) during a feedback planning or plan checking phase. modification of assembly strategies to fit particular geometric environ results in building conditional plans. Then strategies are examined for failures and the planner generates tests (monitoring actions) and corrective actions (which are either conditional branches, or reactive e.g. if the robot manipulator is on the verge of overturning a workpie pushing it with a peg, then retract the hand a little to stabilize the si and then replan the action). If the plan contains many such reactions many contingencies for the same situation, the agent may become too s respond to some of the most time-critical of these contingencies. The s is to use the framework developed here to choose among these conting Further refinements of the plan-skeleton paradigm include symb computations of the effects of uncertainties [Brooks, 1982] to identify an the most significant ones by making inferences about uncertainties and them in computations, as well as using formal program proving techniq deal with these uncertainties [Pertin-Troccaz and Puget, 1987]. All discussion shows that, even if the robot manipulator programming dom not, as a whole, a high-level knowledge intensive domain (in the sense before), the formalism presented here can still be applied if the s uncertainty-related contingencies becomes too large and if their tre requires conscious actions (as opposed to just locomotive reflexes).

Besides the domain characteristics, the agent's capabilities are important in this discussion. If we have an ideal agent with unli resources and unlimited speed of computation, then the entire formalis become useless, even in the real-world, high-level domains presented l However, if we are again interested in the real world, then it is only nat assume that the agent has limited resources and that the number of contingencies for which it has to prepare exceed both its conditional p capabilities, and its real-time execution capabilities. In such cases, the exerts time pressure on the agent's limited resources. Therefore, the needs to be able to decide which contingencies to prepare treatment f which to ignore at planning time. These are the types of agents and d for which the framework developed here is useful.

2.4. Related Work

We make here a brief review of other work that is relevant to problem of how to combine planning and reaction to achieve the performance of the agent in a particular environment. The purpose section is to place our work in the global context of related research outline its original contributions.

Planning (describing a set of actions expected to allow the agen achieve a given goal) has been a central problem in AI since its beginnings [McCarthy, 1958]. The techniques proposed have evo considerably, and so have the application domains. We classify t techniques into several classes, according to the ways they combine th fundamental control modes described before: conditional planning (also called here classical planning or simply planning) and reactive planning (also simply called reaction). These classes are:

(i) purely conditional planning techniques
(ii) purely reactive techniques
(iii) static combinations of planning and reaction
(iv) techniques to shift from planning to reaction
(v) techniques to decide at execution time whether to (re)act o continue the replanning process
(vi) techniques to decide at planning time which contingencies prepare reactions for

A lot of early planning work has been conducted towards specif robust techniques for conditional planning. The systems produced (e.g. [Fikes and Nilsson, 1971], NOAH [Sacerdoti, 1975], MOLGEN [Stefik, 1981], TWI [Chapman, 1987] to almost randomly name just a very tiny subset si exhaustive summary would be well beyond the scope of this section) wer to solve increasingly complex problems. Although some of them had fa for monitoring their plans execution and responding to some conting (e.g. PLANEX for STRIPS [Fikes & al., 1972]), these facilities were very limi and worked only in well-structured domains, based on the existence of matching the contingency in the original conditional plan. More flex and higher response speed was needed to build systems for real-world t

The need for reactivity to the dynamic aspects of the environmen addressed by building systems which operate on a perception-action without relying on an abstract representation of the environment [I 1991]. Horizontal layer decomposed systems [Brooks, 1986; Kaelbling, included such reactions while still being able to pursue high-level goal their reactions were limited to the types of locomotive, low-level preco reactions which we described earlier and which do not make the object work.

Realizing full reactive behavior (reaction plan planning) has l proposed through universal plans [Schoppers, 1987] which are exha conditional plans, and therefore are prohibitively expensive to produ any reasonably complex domain [Ginsberg, 1989]. Situated Control [Drummond, 1989] are used for situation-based plan indexing, to redu non-deterministic choice in the case of plan nets. They may be used incomplete alternative to universal plans, in those cases when there enough time to build the entire universal plan. An incomplete univers may not contain any answer to a problem, while missing situated contro do not necessary preclude a solution (which may be fou nondeterministically); they only ensure a solution when they are spe This approach maximizes the use of planning time and takes into a planning resource limitations, but without taking into account any ex time limitations of the agent.

Pengi [Agre and Chapman, 1987] is a purely reactive planning sy which uses sensory input to index structures for possible subsequent However, Pengi cannot completely represent most real situations due t uncertainty and the limited information available about other agent processes.

Due to the shortcomings of pure reactive systems, researchers subsequently concentrated on integrating planning with high-level re [Firby, 1987] uses Reactive Action Packages like stored reactive plan integrate planning and reactive responses. However, reactive planni used without time considerations, while we allow the agent to tr dynamically replan its course of action if there is enough time to do i only prepare to react to critical events. [Hendler & Agrawala, 1990] imp reactive planning systems on a guaranteed scheduling, real-time op system using the Dynamic Reaction model: an agent performs an activit either its goals lead it to select some new action, or some event in the forces it to react, thus integrating planning and reaction in a co environment. [Georgeff & Lanski, 1987; Georgeff, 1989] propose architecture (the Procedural Reasoning System) that is both highly r and goal directed. They store (reactive) plans, called Knowledge Area procedural form, supplied in advance. [Cohen & al., 1989] monito execution of the Phoenix agents' plans and use three mechanisms handling unexpected events: low level reflexes to stabilize the situation, recovery and replanning implemented as high level cognitive actions envelopes as a general monitoring mechanism. The agent always prepa the same fixed set of reactions, without considering the characteristics plan or of the situations that might be encountered during its execution systems have limited flexibility since the set of reactions is limited, alwa same, and always available in its entirety to the execution components.

Hardware implementations of reactive plans into agents whose a are guided by overall goals have been proposed in [Nilsson, 1988; Continuous actions are modeled using T-R trees (teleo-reactive, i.e. goal-directed and ever-responsive) to build a reactive program v execution produces circuits to control the agent's actions. Selective re would be very important here because of the various costs associated hardware implementations.

The next step on the research path towards agents with better res performance was to devise techniques which shift some of the sys activities from planning to reaction, with the aim of producing increa reactive agents. [Mitchell, 1990] combines reactive (stimulus-response search-based architectures to control autonomous agents. Explanation learning techniques [Mitchell & al., 1986] are used to extract r (condition-action pairs) from plans to make the Theo-Agent increas reactive by learning plans into reactions: the agent first tries to react, plan. Scaling issues for the approach are briefly mentioned, and a solu proposed based on selective learning invocation using a utility fur similar to the one suggested in [Minton, 1990]. However, as we ment before, there are too many characteristics of the situations and contin as well as of the agent (planning and execution modules) which are not into account by this utility function. This fact is even more important rules are tested in sequence for reaction, which yields a high cost of r at execution. [Martin & Allen, 1990] propose a two-level architec consisting of a strategic planner (generating high-level goal descrip which sends commands to a reactive system which must fill in the details use statistics to constrain the probability that the execution modul accomplish a particular task. Reactive behaviors are learned selectively, statistical estimates on the utility of these actions versus the utility of components. But once learned, the reactions are always available to execution system. Soar [Laird & Rosenbloom, 1990] also provides a combi of reactive execution and planning seen as essential behaviors of autonomous intelligent agent. Plans are learned into reactions chunking, and afterwards all reactive plans learned are always availal the executor. The authors express their concern that after learning too such reactions, the responsiveness of the system may be significantly r but do not attempt to address this problem.

These works concentrate mainly on how to prepare reactive resp and tend to use them in such a way as to substitute regular plannin approach differs from these others in its recognition of the complem strengths and weaknesses of the two modes, and in its full integrati planning and reacting within a single agent. A recurring, unaddr problem in these works is the value (utility) of reaction. While we believ learning such reactions is very useful in real domains, we also believe this utility problem should be addressed at planning time, and not (o learning time. The work described in this thesis is aimed precisely towar goal. In the next chapter, we will define a framework to select only relevant events associated with a given situation. Reactions to them incorporated into stored reactive plans, depending on several factors s event criticality, reaction time allowed and exhibited, load of the a reasoning capabilities and other resources, and reactive plan size, as w on the desired behavior pattern for the agent. Our main problem is to which contingencies to prepare reactive responses for, in each situation is in contrast with most of the research cited above, where the autho concerned mainly with improving their systems' performance by tryi react (and maybe cache) the most frequently used plans. Our selection will necessarily be much more complex than the utility measures propo far.

However, the utility of reacting versus planning can also be, and lately already been, addressed at execution time. [Horvitz, 1989] deve decision theoretic framework to reason about the value of continui reflect about a problem vs. taking an action to try to solve it, at executio using the expected value of computation (EVC) as fundamental measur attempts to optimize behavior under resource constraints by inte reaction with deliberative reasoning (replanning). However, he ignore overhead of retrieval of a reaction and the computation time while taki account only limited other resource constraints (e.g. memory cost) whic not be the most relevant ones for real-world agents. He also assume reactions are always available and only attempts to decide, at executior whether to react or to replan, and is not concerned with such decisi planning time (clearly, some contingencies do not allow time for metalevel deliberations at run-time, before taking an action to respo them). [Yamada, 1992] uses the notion of success probability to determi best time until which dynamic replanning may continue and when ex of the action should actually start. Again, the computation is done at ex time.

The sixth category of techniques which we have identified at beginning of this section involves methods to decide, at planning tim which contingencies to select for preparation of reactive responses i plan, and which to ignore and leave for dynamic replanning at executic if such a contingency will arise. The problem is occasionally mentioned literature, but without being analyzed in detail and especially wi proposing any solutions to it. While discussing the CIRCA system, [Muslin al., 1994] make the most comprehensive presentation of the problem t were able to find. They recognize the limitations that exist in the execution resources, and attempt to divide the main plan into smaller and create reactive plans that guarantee the achievement of critical However, there is no analysis of how to partition the set of goals guaranteed and unguaranteed ones (when the system cannot gua responses to all of them). CIRCA only tries to build guaranteed plans by into account only the time allowed to respond to a contingency. contingency characteristics relevant for the decision process (like crit and probability) are mentioned as necessary to be considered in future but they are not actually used here. Control level goals are linked t system's safety, which is not always necessary (in our work, any chan the environment that was not expected as a result of executing the mai is considered a contingency). CIRCA also partitions the goals into just subsets according to a system designer specified priority: critical or not.

We are unaware of any previous research towards a solution to general problem of deciding whether to prepare a reactive response contingency or not; therefore, it is here where the work described i thesis has been concentrated.

As shown before most research to date is concerned either employing only one of the planning or reacting control modes, or attempts to turn a system to become increasingly reactive and rely as l possible on planning. All the reactive responses are always available t agent executing a plan, and they usually tend to take precedence ov (re)planning alternative. This approach can only work in either very task environments, or for idealized, unlimited resource agents. In our we take into account the real-world constraint of limited resources for that have to act in stressful, resource-demanding, real-time situatio which reaction does not come for free. Therefore, we assume that importance of regular planning makes it irreplaceable, but the a performance can be significantly improved by selectively preparing r responses only for those contingencies that are critically enough to them. We work towards integrating planning with reaction, instead o enabling the agents to shift from planning to reaction. [Hayes-Roth, proposes a paradigm for integrating planning and reaction u opportunistic control of action: run-time control conditions trigger a su possible actions, strategic plans constrain intended actions, and the between possible actions and strategic plans controls action execution.

Other work, directly related to various subsections of the thesis, briefly surveyed when relevant.

Chapter 3
The Approach

In this chapter we describe our framework for deciding, at plar time, whether to prepare a reaction for a given contingency in a c situation. We first define a few terms which we will frequently use:

- a *plan* (*conditional plan*, or *main plan*, or *conventional plan*) is a (possibly conditional) time dependent, partially ordered set of a and expectations (figures 2.1 and 3.1.a).
- an *action* is the application of an operator to the current state. It y new state, which may be identical or not to an expected state.
- a *contingency* is any state of the world entered by the executing while following a plan, which is not: (i) a direct consequence executing the actions of the plan up to that point, or (ii) an exoge generated state of the world assumed in the design of the p Therefore, a contingency does not necessarily affect the agent or plan execution, and when a contingency does affect the plan, it is necessary that it will negatively affect it. For example, a conting may be a state which is not the current expected state according t plan execution, but is a state which should have been reached alon way, after executing some more steps of the plan. The agent may c it and use it to skip the unnecessary steps in the plan, for examp the same way as it was done with triangle tables in [Nilsson, 1984] simplify the exposition, we also use the term *contingency* to mear event, fact or sign that was not expected as a result of the p execution, and which triggers an (undesired) change in the state o world, not expected at that time in the plan, i.e. which characteri state as a contingency according to the previous definition.

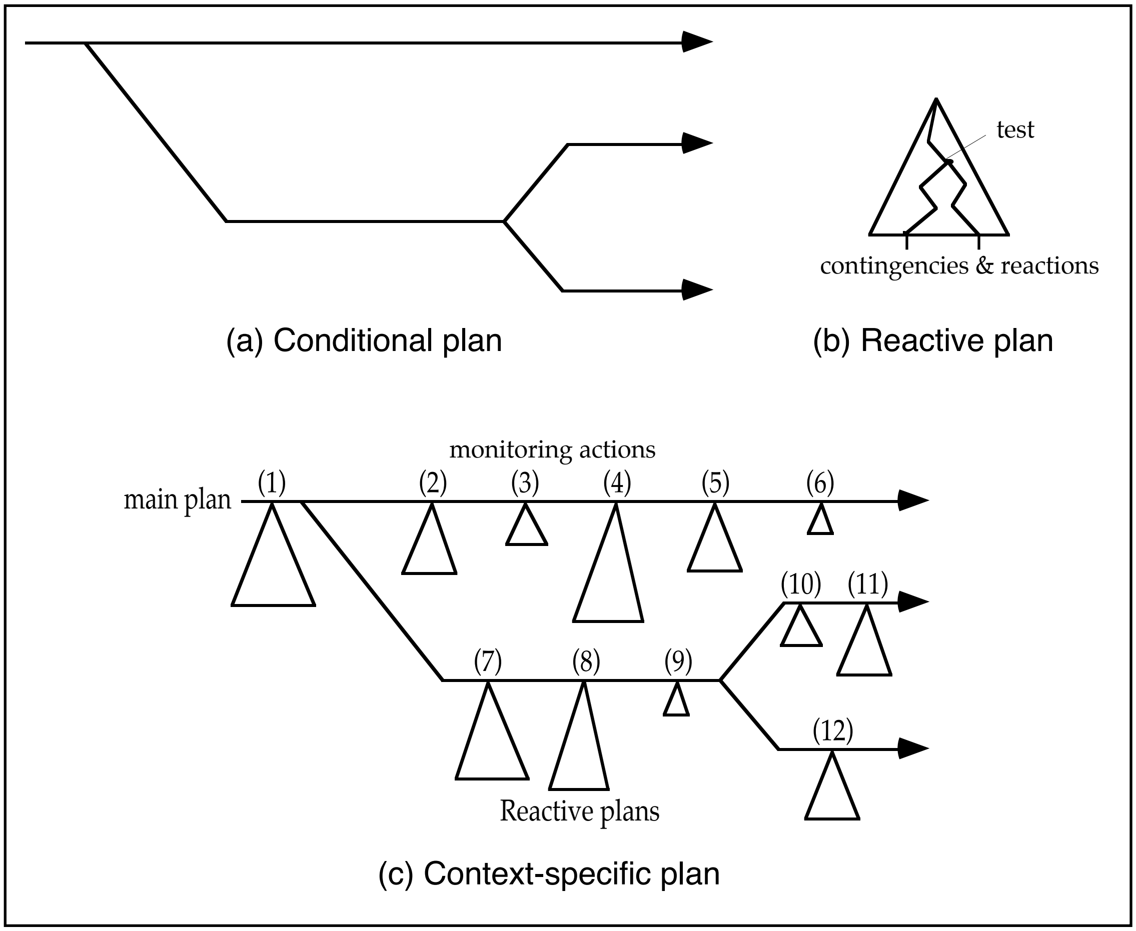

Figure 3.1. Types of plans

- a reaction is a perception-action rule of behavior, usually stored computationally efficient form. The action part may be a short sec of actions which are enough to either solve the problem generated contingency, or at least to extend its deadline long enough to allo replanning of the entire solution under the new circumstances.
- a condition is a pair contingency-reaction; there may be more thar reaction which can solve the same contingency, and there may be than one contingency which can be solved by the same reaction.
- a reactive plan is a set of tests and reactions (possibly arrange hierarchically for efficiency reasons [Ash & Hayes-Roth, 1993] therefore represented as triangles in figure 3.1.b) able to solve an of a set of contingencies.
- a context-specific plan is obtained from a conditional plan by augmenting it with monitoring actions and reactive plans for ce contingencies (figure 3.1.c). It deals with these contingencies in a and usually incomplete way, as opposed to the conditional plan prepares in advance for a full treatment of the possible situation were taken into account.

The basic approach to obtain a final context-specific plan for a problem starts with a conditional plan (produced by a conventional p to achieve the main goal of the problem. The agent has a knowledge b contingencies that may appear during the execution of plans, togethe proper reactions to them. After developing a plan, this knowledge is u analyze it and to identify situations of interest, that is, those points in t for which the agent knows of possible contingencies and how to respo them.

The general agent architecture to do this is briefly discussec appendix 1. In the rest of the thesis, we assume that the agent has a decided upon such a situation and has identified the set of conting which may be associated with it together with their appropriate re responses. Now the task of the agent is to decide for which of t contingencies to actually include responses in a reactive plan which subsequently be attached to the main plan at the appropriate place (s by the particular situation isolated before). The context-specific plan i completed by augmenting the initial main plan with monitoring actior reactive plans for the critical contingencies (figure 3.1.c). Monitoring a can be attached to the plan even if reactions to their contingencies a (e.g. when the contingency is important enough to be watched for, but its likelihood of occurrence is low enough, or the time allowed to respon is long enough for replanning).

In the next section, we first analyze a simple problem and try formulate an intuitive solution. We then formalize this intuitive solut the rest of the chapter.

3.1. Intuitive Solution

Let us revisit the driving problem presented in the previous cha and attempt to analyze it in more detail.

In section 2.1 we formulated the problem of an agent which comi every morning by car from home to work, and at some point A along the passes in front of a school while driving straight, at 25 mph. The coı takes place at a time when children are at school, or go to school. The knows its route well enough to know about a few contingencies that may while on this portion of its route. Table 3.1 lists a partial set of contingencies, and the best reaction for each of them known to the Notice that the contingencies are dependent on the characteristics c actual situation described. Here are some of these dependencies: contingencies depend on the type of plan used (e.g. if the agent uses transportation, than it need not be concerned with hitting a child, sinc not in control of the car), on the action involved (if the current action be driving on a freeway, then the likelihood of having children runn front of the car would be much smaller), on the context of solving the p (if the same action takes place during vacation time, when that scho closed, then again the likelihood of having a child run in front of th decreases a lot), and so on. In the next section, we rigorously define the of a situation, and then precisely characterize this particular situation example of our definition.

In order to be useful for our purpose, the notion of a situation (a associated characteristics) must be much more rigorously specified. Als contingencies must be expressed in some structured language in order to a better representation and usage (e.g. it is important whether the car slowly or fast, whether the child runs from left to right or from right t and so on). We detail these specification requirements and present forı to facilitate their expression in the next three sections of this chapter the next chapter.

	Contingency	Reaction
1	Child runs from right, 20 m in front of car	Brake hard and steer right
2	Car crosses w/o priority 20 m in front, from right to left	Brake and gently steer right
3	Car in front stops suddenly	Brake hard
4	Cat runs across street, 20 m in front	Brake hard and steer right gently
5	Traffic light changes red 40 m in front	Brake hard
6	Tire explosion	Brake gently and do not steer
7	A deep and medium width hole detected 30 m in front	Brake hard and steer right gently
8	Airplane lands in front of car	Brake moderately hard
9	Brake malfunction light turns on	Brake gently
10	Engine overheat light turns on	Brake gently to stop the car
11	Loud radio turns on suddenly	Adjust radio volume
12	Meteor falls on the trunk of the car	Accelerate hard

5

13
A ball pops in the street, from the right, at 20 m in front
Brake hard and steer right

Table 3.1. Set of contingencies for the car driving domain

Our problem is to decide which of these contingencies are cri enough to require the agent to prepare in advance reactive respons them and which should be ignored at planning time. The solution ha phases. In the first phase, the agent must order the contingencies accor the value of reacting to them; then taking into account the characteris the planner and the limitations of the agent's run-time resources, it mu out how many (and actually which) of the contingencies can be taken account for reactive treatment. In order to be able to define the val reaction to a contingency and to be then able to order the conting according to this value, we have to identify the characteristics contingencies which influence this reaction value. These characteristic defined not for a contingency alone, but for a condition (contingency-response) in a given situation (as seen above, t characteristics can vary from one situation to another).

One characteristic which has been recognized by earlier research remarked in section 2.4) is the likelihood of appearance of the continge that situation. We have already discussed how the same contingency ma different likelihood in different situations. Also, different contingencie have different likelihood in the same situation. For example, in our c child running into the street is less likely than encountering a red light, but more likely than having a plane land on the street in front car.

Since reactive response is geared especially towards satisfy real-time deadlines, of special concern is the time pressure exerted b contingency upon the agent. This time pressure (or urgency) is inv proportional to the actual real time allowed for the agent to act in resp the contingency. Clearly, responding to the child contingency is more than taking care of the radio which has just turned on by itself. On th hand, the child running into the street and the ball popping up in fron car at the same distance, allow for the same time of response, i.e. exe same time pressure onto the agent.

But the value of reacting to a contingency is also determined by gravity of the consequences presented by the contingency if no acti taken in the allowed response time. Obviously, the consequences are more dramatic in the case of hitting a child, than if the car hits a ball.

And finally, there is one more characteristic of the conditions tha to be taken into account. This characteristic is more closely related response associated to the contingency, and it takes into account the side-effects that may be incurred if the reaction to the contingency is t time. For example, the side-effects of avoiding the child by braking har possibility to be hit by the car following our agent's car) and steering (the agent's car may hit the sidewalk, or a pole on the sidewalk) are th as for avoiding the ball through the same maneuver, and can be signi higher than the side-effects of adjusting the radio.

We assume that the agent's knowledge base contains, along with contingency and reaction, a set of values for these characteristics (the be obtained from experts in the domain - as we have done it, or th automatic learning methods). These characteristics have different weig deciding upon the value of reacting to a given contingency. As we shal these weights are not fixed, but they are dependent on the application and also on the behavior model according to which the agent acts. We sh now restrict our discussion to a generally accepted (by the experts i domain) "normal" behavior, and will briefly discuss other types of beh in section 6.3. Under this behavior model, the highest weight is associa the time pressure characteristic, followed by consequences and likelihood. However, if the side-effects are much higher than consequences, then the agent is probably better off by ignoring contingency at planning time.

Therefore, a driving agent will give highest priority to the c running into the street contingency (since the time pressure is very and the consequences are also very high), and will give a very low prio the ball contingency, since the side-effects of doing a dangerous man outweigh by far the consequences of hitting the ball. The traffic light t red contingency will follow the child one, followed in turn by the ai landing and the loud radio turning on (since both have low likelihood, airplane has much higher consequences and time pressure). contingencies listed in table 3.1 are actually ordered according to the behavior model described by a panel of experts whom we have interv (section 6.1 presents more details about our knowledge acquisition proc this domain). At first glance it may be surprising, for example, that th contingency was placed after the radio contingency; remember howeve we are only interested here in preparing reactions for these conting Therefore, this ordering says that, if the agent has enough resources, try to prepare a reaction to the radio contingency (although the va reacting to it will be pretty low), but should avoid as much as possib prepare a reaction to the ball contingency, since the side-effects of reac it may be much higher than the consequences of not reacting equivalently, the benefits of reacting).

The second phase of our solution involves deciding which of t contingencies will actually be included in the reactive plan, by taking account the characteristics of the reactive planner and the limitations agent's resources. The characteristics of the reactive planner (specified reactive planner model) allow the agent to estimate the complexity of is the contingency and its reaction from the reactive plan prepared fo entire set of selected contingencies associated with that situation. complexity is direct proportional to the time needed by the agent fro moment it detects the existence of a contingency and until it can s reaction to it. However, this time is further influenced (i.e. increased) availability and limitations of the agent's resources, specified by an model (e.g. computational overhead). For each contingency included i reactive plan, this response time has to be smaller than the time allow the contingency before the (re)action has to be taken (otherwise the r to that contingency becomes useless). Therefore, given the reactive p model and the agent model, we have to analyze each contingency asso with the situation, in the order specified by the first phase of our anal our example, we will always include in our reactive plan a response t child contingency, since it has the top priority. We will also include i plan a response to the car crossing contingency, if we estimate that the will have the resources to react to both contingencies in time, and so on. reach a contingency which cannot be responded to in the allowed time while still being able to respond to all the contingencies included i reactive plan before it, then this contingency will be left out. However process continues until all contingencies have been examined, since contingency further down the list may allow a longer response time, still allowing time to respond to all the already included contingencie example, assume we have time to respond to only two contingencies wit high time pressure, and to some other contingency with much lower pressure. Then we will want to include the child and car cros contingencies (which are the first two on our ordered list), ignore th stopping and cat crossing contingencies for which we do not have ti respond, and include the red traffic light contingency which follows i list, because it allows for a much longer response time. Such a policy (wh rigorously defined in section 3.4) makes optimal use of the agent's exe time resources, as justified in chapter 5).

In the following three sections we define our framework, along lines of the intuitive analysis presented here, and in chapter 5 we m brief analysis of some of the theoretical properties of this framewor chapter 6 subsequently then present a few more examples of applyin framework in other domains like anesthesia and intensive care monitoi

3.2. Framework for Reaction Decision

In the following sections we define our framework, along the line the intuitive analysis presented above. We specify a consistent framew help decide whether the agent should prepare in advance to react to possible contingencies, or whether it can ignore them at planning tim can replan at execution time to deal with them. As seen before, the inclu monitoring actions and/or reactive responses for a particular continge a plan may depend on a large number of characteristics of the envir the contingency and its response, and on the relations between them, as on the models of the different factors involved in this process: the the agent and the reactive planner. They also depend on the set of contingencies possible in the same situation (how many, how critical, an complex their reactions are) vs. the agent's capabilities. To help visuali heuristic rules that take these decisions, we define a few multi-dimer spaces and the relationships among them. The position of a continge these spaces determines whether or not the agent reacts to the event.

3.2.1. Overview of the Framework

We begin with a general presentation of the interactions among components of our framework, and in the subsequent sections we pres detail each of these components.

Figure 3.2 presents a schematic overview of the framework desc here. The entire framework is used to decide, for a given condition contingency-reaction), whether the agent should include the reaction contingency in the reactive plan which is prepared for the situation consideration. Therefore, given the condition and the situation, framework has to provide the means to associate a criticality value t contingency. This criticality reflects the value of reacting to the conti (using its associated reaction, if it appears in this situation), as oppo leaving the agent unprepared to respond to this contingency and hopi it will be able to solve it by dynamic replanning if the need will arise. reaction value is high enough, the agent will at least monitor for occurrence of this contingency during execution of this phase of the However, the agent may not be able to prepare for all contingencies criticality high enough to be monitored for.

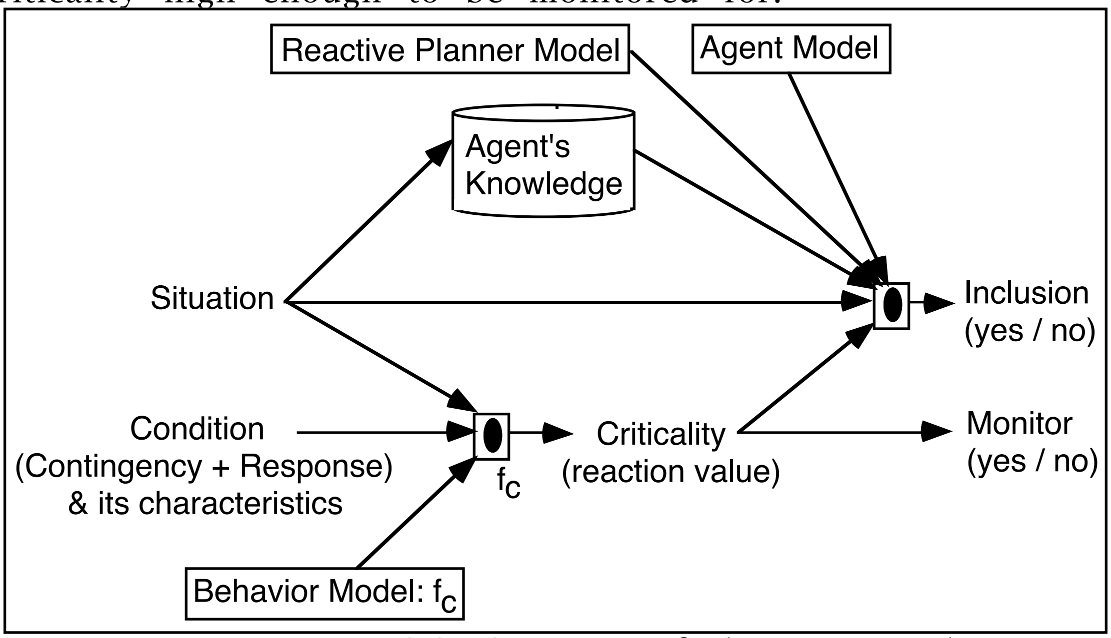

Figure 3.2. Overview of the Framework

The decision of whether to include the reaction to this contingen the reactive plan is taken based on the characteristics of the situatio time pressure exerted by the contingency upon the agent (or equivaler time allowed for response by the contingency), and of course the critica the contingency, compared with the criticalities of the other conting known to the agent to possibly appear in the current situation. The cr values induce an order relation on the set of contingencies associated situation, and the agent first attempts to include the most critical of contingencies for reactive response. All the contingencies (taken fror agent's knowledge base) associated with the current situation are con in turn for inclusion, in the order of their criticality value. When re the stage where the current contingency is analyzed, all the conting applicable in the current situation, with higher criticality, have been analyzed, and for some of them (not necessarily all) the agent has deci include reactive responses in the reactive plan. The current contingen be included in the reactive plan only if the agent using this new reacti will be able, at execution time, to respond to this contingency in its a time, while still being able to respond in their allowed times to all contingencies already included in the reactive plan. In order to tak decision, our framework needs a model of the characteristics of the r plan built by the agent, as well as a model of the execution time charact of the agent resources and their limitations.

Figure 3.3 presents in more detail the source and flow of inform through our framework. Each situation has a number of characteristics, therefore represented as a point in a situation space. This represe allows for flexible generalizations and for the representation of sets of situations as regions in the situation space. Similarly, the characteristic contingency will define the dimensions of a criticality space, in which point represents the value of reacting to that type of contingency. Th space used represents the reactive plan characteristics, in terms o resources required by the execution of the reactive plan (given by reactive planner model) and the resources available for execution b agent. The agent model gives indications on how these resources are m by the agent and how they are used by other modules of the agent, as the limitations on the agent resources, and is therefore used in the fin of the decision process. The expert model is used by the framework to i the values suggested by the expert for the characteristics of contingencies, and specifies a set of threshold values for these characte Finally, the behavior model defines the function which computes criticality value for each contingency. Different behavior models ass different values for the same reaction to the same contingency, accord the individual values of its criticality space characteristics. The two c stages of the framework are establishing the criticality or reaction val the contingency, and making the decision of whether to include its r into the reactive plan built for the current situation.

In the remaining subsections of this section we discuss in detail ea the three spaces mentioned above, and then we present a complete sum the entire framework. The following two sections will then describe th critical points of the framework mentioned above.

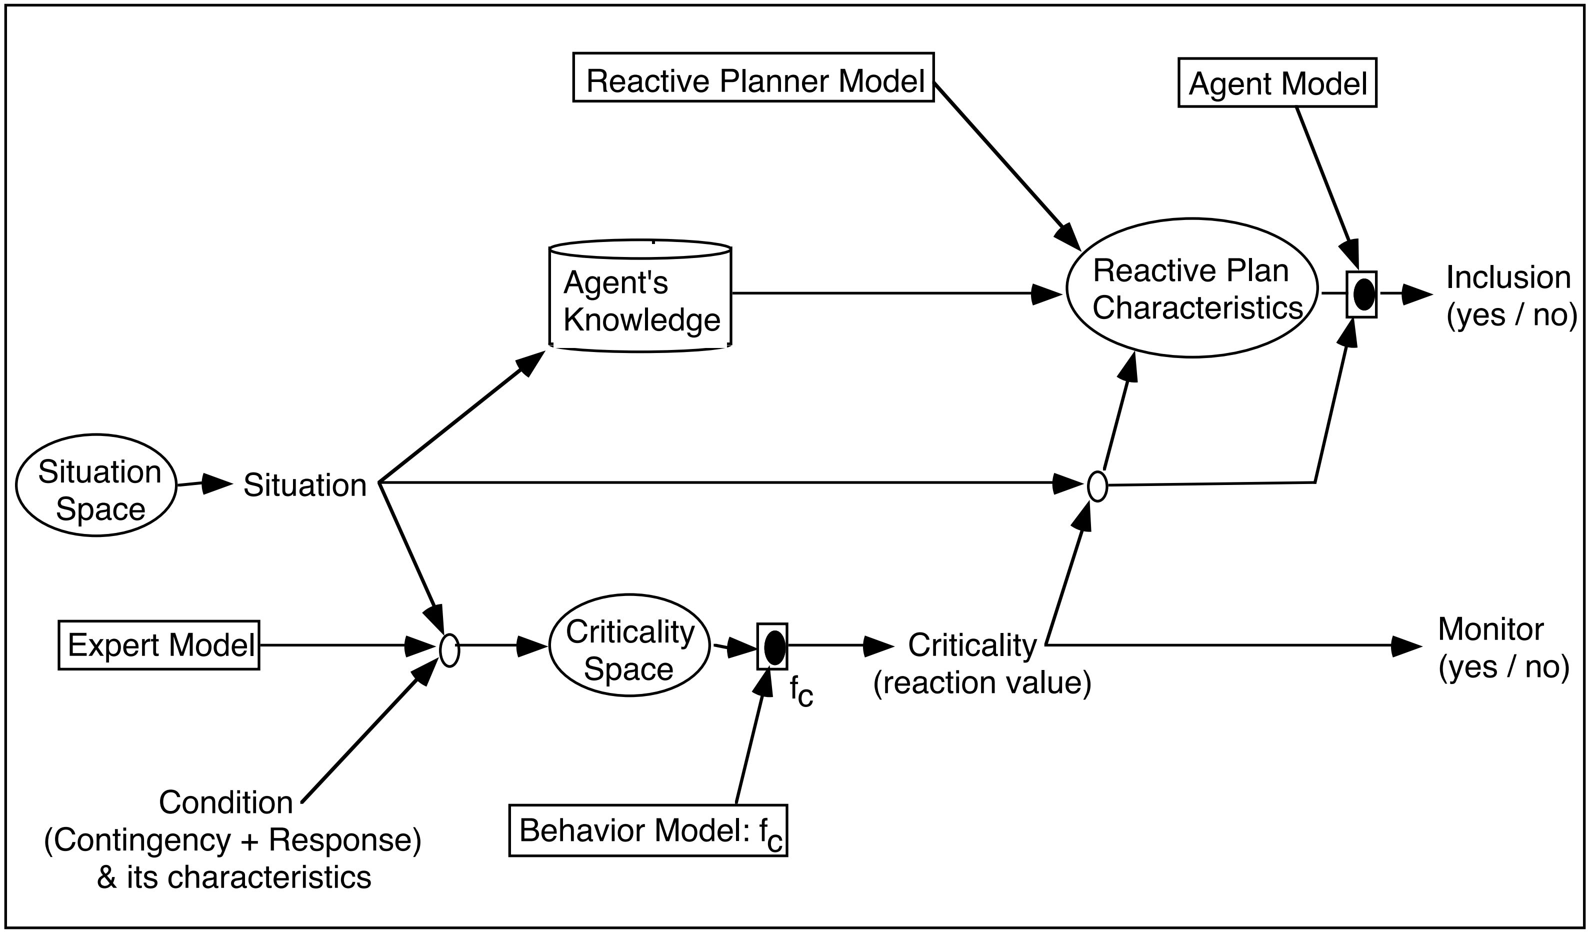

Figure 3.3. The General Framework

3.2.2. The Situation Space

The situation space is the set of all possible situations. Its dimensior the aforementioned characteristics of a situation. A point in this characterizes a general, contingency-independent environment situati state. Situations will be used to index contingency-response pairs ir agent's knowledge base, according to the relevant situation characteris which they may apply. We will elaborate more on the same driving ex used before, and will try to specify it more accurately from the perspec our problem. The seven dimensions of this situation space are:

- problem - is the main problem to be solved by the agent. It is a syr of the problem characteristics and how they can determine the ; situation. An example of problem is to carry a small package of b from home to work. We shall use this example throughout this secti small change in the problem statement can have important influ on the set of contingencies that can be expected. For example, if problem is instead: carry a small package of radioactive material home to work, then an entire subset of contingencies generated b fact that the package contains radioactive materials has to be taker account.
- plan - is a synthesis of the characteristics of the type of main plan to solve the problem. The type of plan chosen by the convent planner is obviously dependent on the problem to be solved. example, the plan may differ depending on the size of the package carried, on its weight or on its content, as well as on the distance traveled. However, even for the same given problem there may l large number of solutions (plans to solve it), and each of them create different conditions with which contingencies may be assoc For example, for our problem, one can choose to walk or to use a r of transportation, and further, to drive or to use public transpoı and further to drive a car or a bike, or any combination of these, a on. Let us assume the planner's choice was to drive a car.
- context - is a synthesis of the characteristics of the environmer which the plan is to be executed to solve the problem. It covers al general aspects of the domain which are not covered by the pre two dimensions. For the driving example, it includes the time of th (it may make a considerable difference for the types of contingenc be expected, whether it is day or night), the time of the year (in v the road is usually more slippery, but the engine is less likely overheat), weather conditions, the abilities of the driver, and sc Suppose in our example the context is a working day morning d the month of May. This means that children are going to school, therefore children and balls can be very well expected into the around the school.
- action - is the action to be currently executed by the agent accord the plan. Since the contingency preparation process is an off analysis of the main plan, "current" here means the currently an time point of the plan. Non-execution of planned actions (mis actions) may also be represented on this dimension, since conting may occur both associated with the execution of actions in the plan (e.g. steering to the right may cause the car to slip sideway well as with non-execution of an action (e.g. not steering to the

when the road turns right may have obvious consequences...). In example the action is just to drive straight ahead on street S at a sp 25 mph.

- expectations - are descriptions of situations (changes in the state environment) along the plan path. In order to monitor the executi the plan, the agent looks for some important such states which prespecified at planning time. We call these states milestones. achievement (or not) of a milestone may determine the agent to c the conditional plan branch which it is following, and therefor change the set of contingencies for which it is on the looko According to the way they may be generated, there are two kind expectations which must be taken into account when definin situation:
  - internal expectations - due to actions performed by our agent while executing the plan (e.g.: an attained milestone may be entering freeway, as expected, while to the contrary, an unattained mile may be a situation in which the agent did not enter the free although this was expected as a result of executing a set of steps). Such an occurring state change can be foreseen, and if change does not occur, it becomes a contingency: it may signal something went wrong with the plan execution, and therefore agent should try to find out what and replan, but in the meanti should be on the lookout for a certain set of contingencies that also appear in this situation. For example, due to driving on stre the agent expects (as milestone) to arrive in front of a school. does not, then maybe the plan was not entirely correct and the is somewhere else than it should be at that time. It should the react (attempt to stop) and replan: attempt first to find out wher (e.g. by reading the street signs), and then replan its route there on.
  - external expectations - due to other independent agents which wo the same environment (e.g. changes in traffic lights). These ag may generate contingencies by themselves, since they activ change the environment; their actions may have a certain non degree of correlation with the actions of our agent, or may be t uncorrelated. For example, the traffic light is an agent whose ac may be somewhat correlated with our agent's actions if our a approaches the traffic light from some direction where there street sensors or other traffic lights synchronized with this otherwise, the traffic light's actions are totally uncorrelated the actions of our agent. Two kinds of events may be distingui here too: (i) something may happen (like the signal change) or something expected may not happen (e.g. a malfunctioning signal which does not change after a long waiting time period) the example situation we have been building in this section possible external expectation might be to notice children in the (since it is a working day morning in May and we are in front school). However, this is not a milestone: it is possible that children may be in class at that time, and this fact does not alt any way the execution of our main plan.

m time - this basic characteristic of planning problems will appear in of the abstract spaces we consider, although with different mea (when the possibility of confusion arises, we will denote the t dimension for the situation space with $time_S$). Here it represents the amount of time elapsed since some action was taken or since a situ change was noticed, or the amount of time allowed until a situ change must appear. It is therefore strongly coupled with expectations dimensions (expectations become more or less stro with time passage). For example, if we allow for 3 minutes from moment we start driving on street S until reaching the school an expectation is not met, then something wrong may be going on (e traffic jam, or a deviation from the route) and the agent should replan (or maybe first to react and then to replan) for an alte route.

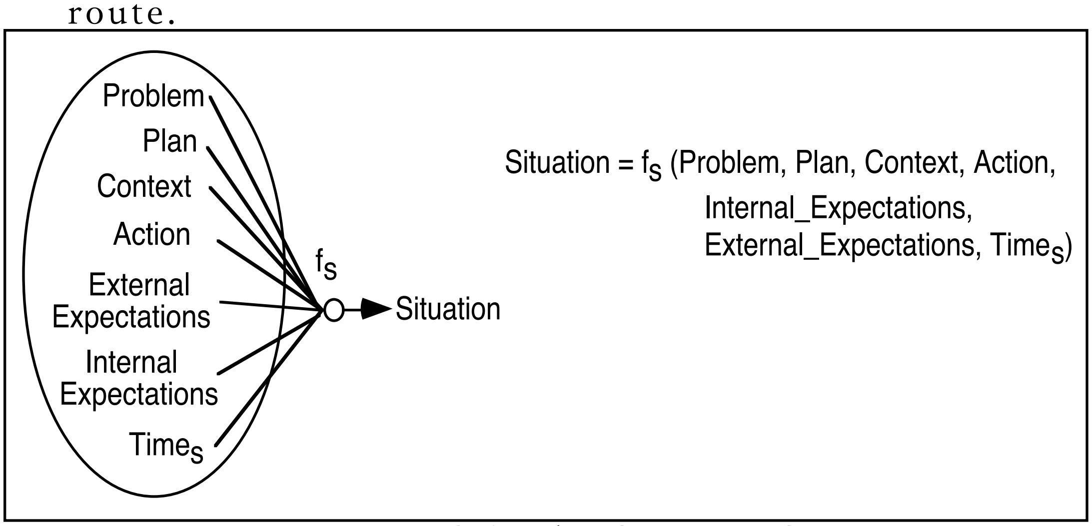

Figure 3.4. The Situation Space

The values along each dimension of the situation space are descri of those dimensions, as given in the example built during this sectio summarized in section 3.2.5. A point (called situation) of this space, fully defines (for our purposes) the agent's situation, that is: the action execu the current expectations in the course of executing a certain type of p solve a given problem in a specific general context or environment. W use it further to determine whether the agent should prepare or not a for a contingency "in the current situation". In chapter 4 we pres representation formalism for the values of the situation space dime which allows us to group situations into classes to facilitate the stora knowledge and the reasoning and knowledge acquisition processes fo agent using our framework. Figure 3.4 summarizes the functi dependencies described here.

With each point in the situation space, there is a (possibly null) s contingencies (and responses) associated (known to the agent throu knowledge base) for which the agent has to further decide whether to for and to prepare reactions for. Let us suppose that the contingencies by our agent to be associated with the situation described in this secti the ones listed in table 3.1. However, we shall mainly discuss and compa characteristics of only two of these contingencies, which have essential same reaction: (i) children running in the street in front of the car, and ball appearing in front of the car. As the need will arise, we will refer tc contingencies in the set for comparisons too.

3.2.3. The Criticality Space

The criticality space describes the characteristics of a contingency its associated reaction in a specific situation, and helps in establishir value of performing the reaction when the contingency appears in situation. In the previous subsection we used the situation space to eva situation, independently of the contingencies that might appear in it. H evaluate the criticality of a contingency, dependent on the situation in it occurs, but independent of the set of other possible contingencies f same situation, and independent of the characteristics of the reactive and those of the agent. Resuming our driving example, we continu exemplify our presentation by analyzing the two contingencies asso with the situation described during subsection 3.2.2. The four dimension situation-dependent values) defining the criticality space are (figure 3.:

- time - is the time deadline, or the urgency to correct the problem by the contingency. This is in contrast with the time dimension fo situation space introduced in the previous subsection, wh represented the time allowed to pass until a contingency is declare actually use two strongly correlated values here:
  - $Time_{rc}$ - is the actual real-time interval allowed to pass (witl consequences) between the time a contingency is detected and the corrective action is taken.
  - $Time_{p}$ - is the corresponding time pressure acting upon the agent inversely proportional to the real time (the proportionality fact a parameter of the expert model).

  In our example, in both the child and the ball case, this is the dy planning time available before the action must be taken in ord avoid collision, from the moment the contingency is detected. This is shorter than, for example, the time allowed to respond to the turning itself suddenly loud. Therefore, the time pressure is r higher in the first two cases than in the radio contingency.
- consequences - is a summary of the gravity of the consequences may appear if no action is taken (before the time deadline) in res to the contingency. This value can (but need not) be situa dependent. In our example, hitting a child can be fatal, and this will be very high. But hitting a ball is usually no big deal, so its will be small.
- side-effects - is a summary of the gravity of the consequences that occur as a result of reacting, and therefore this characteristic is r dependent on the reaction and the situation, and less dependent actual contingency. Alternatively, it is a measure of the risk of being able to reach the final goal anymore, once the reactior executed. In our case, in order to avoid hitting the child or the when driving a car, the same reaction is indicated. It is a dange maneuver (braking hard implies the possibility to be hit by the following our agent's car, and steering right implies the possibility the agent's car may hit the sidewalk or a pole on the sidewalk) an yields a high value for the side-effects characteristic in this case, significantly higher than, say, the side-effects of adjusting the rad

- likelihood - this dimension summarizes the probability of occurren a given contingency in a given situation. However, it is importan note that it need not be the actual probability, or not even per correlated to it. It can simply be a value that is approximately corr to the actual probability, in that the relative values of the probab of different contingencies are reflected in their relative likelil values. Initially, this value can be determined from previously k cases in the literature describing the domain, from the estimates expert, or from a theoretical analysis when a sufficiently strong d theory exists. Later on during its lifetime, the agent may adjus according to its own experience. Assuming the agent has no p experience in our example, we initialize the likelihood as medium both a child and a ball appearing in front of the car passing in fr a school, with the likelihood for the ball contingency a little hi than for the child one. They are both higher than the likelihood to an airplane land on the street, but lower than the likelihood encounter a red traffic sign.

The values along the consequences, side-effects and likeliho dimensions of the criticality space are reals in the interval [0,10]. The for the time pressure dimension are real numbers greater than 0; the limit for the time pressure depends on the threshold values imposed expert model, which will be discussed in section 3.3.1. All the values for criticality space dimensions may be specified qualitatively (e.g. for consequences dimension using {very small, small, medium, high, very and are then translated into numeric values. These values are sit dependent; they may be different for the same contingency associated different points in the situation space. For example, the side-effects c proposed dangerous maneuver to avoid a collision with a child or a b much smaller if driving in an empty, large parking lot, than when driv a busy street. The values for the criticality space dimensions for condition and situation must be specified in the agent's knowledge base important to note here that these values need not be very precise in a values. It is enough if they are in the correct order and approximate correct relative values. This is because the method for computing criticality value (section 3.3.2) and the way this value is used further framework are robust (i.e. noise tolerant), making the entire fram robust. We shall substantiate these remarks in chapter 6, when we discuss the experiments we have conducted. Given these relaxed pre requirements, the experts with whom we have worked on the know acquisition part of our experiments were able to specify quickly and wit effort suitable values for the characteristics of the contingencies in experiments.

A point in the criticality space presented here defines an expe value for the reaction to a contingency, versus a dynamically repl response, as shown in section 3.3.2. The agent attaches to the plan s reaction only if the contingency is critical enough with respect to the contingencies possible in this situation, and only if it will have en resources at execution time to respond in time to this contingency as wel all the previously accepted contingencies. That is, as we shall see in s 3.4, not all such reactions can be included, but monitoring actions for al contingencies found to be critical enough (according to an expert d threshold) after this analysis will be included in the plan.

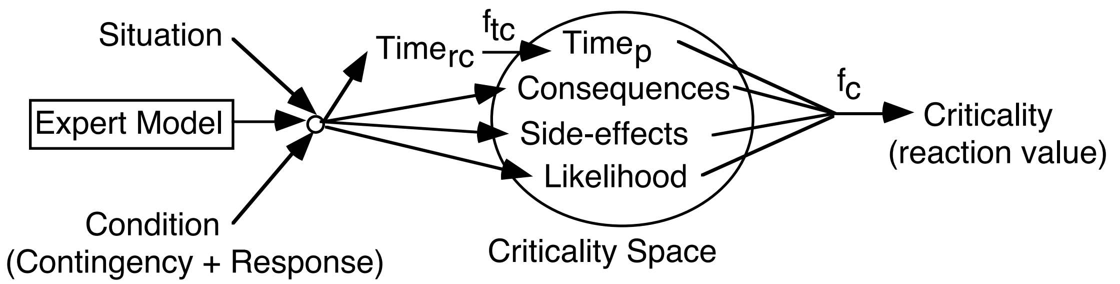

$Time_{rc} = f_1$ (Situation, Condition)

$Consequences = f_2$ (Situation, Condition)

$Side\text{-}effects = f_3$ (Situation, Condition)

$Likelihood = f_4$ (Situation, Condition)

$Time_p = f_{tc}(Time_{rc}) = k / Time_{rc}$

$Monitor = f_m$ (Criticality)

$Criticality = f_c$ ($Time_p$, Consequences, Side-effects, Likelihood)

Figure 3.5 The Criticality Space

Figure 3.5 summarizes the characteristics of the criticality sp defined above, and their relationships (functions) to other elements framework. Functions f_1 to f_4 are implicitly contained in the expert r they are not explicitly used in the framework, since the values for th dimensions of the criticality space are acquired directly form the ex However, for well-structured domains, it is possible that a strong do theory might exist which can explicitly specify these functions.

3.2.4. Reactive Plan Space

The reactive plan characteristics represent one more set of featur consider in deciding whether to prepare a reaction to a contingency We define a reactive plan characteristics space to help us study relationships between replanning a response, versus reacting to the contingency in the same situation. The factors to be taken into accour are the availability of computational and non-computational resources agent, expressed through the reactive planner model and the agent (subsections 3.4.1 and 3.4.2). Here, the values of the dimensions in this will be based on all the elements of our framework: situation, contin criticality, and reactive planner and agent models. Thus, we have bui framework hierarchically, the coordinates of each space of the fram being defined in terms of the values of elements in (and the dimensio the previous spaces.

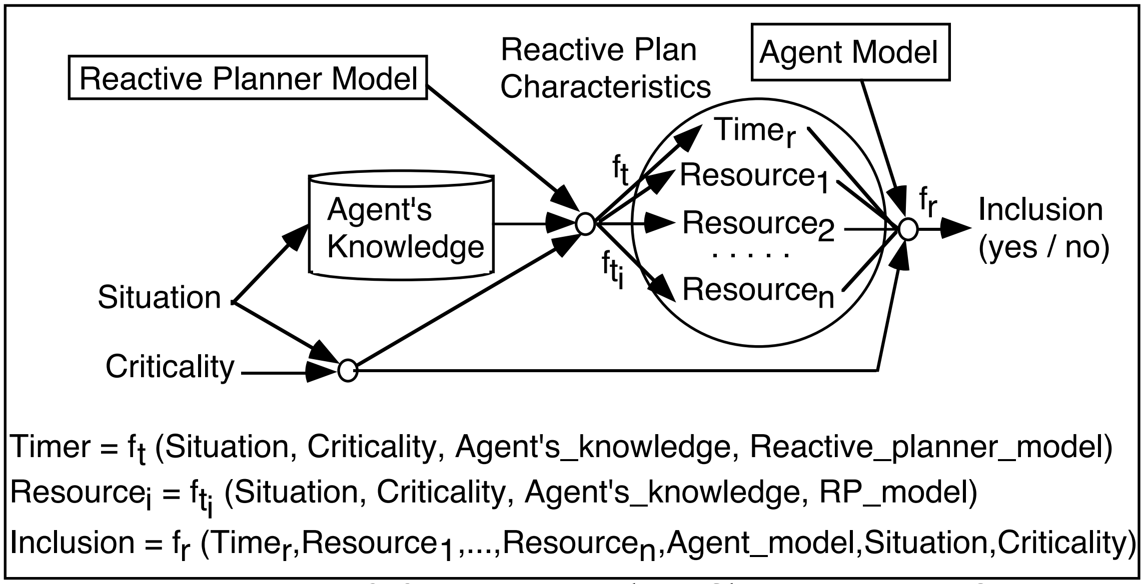

Figure 3.6. Reactive Plan Characteristics Space

The dimensions of the reactive plan space, which also represent characteristics of reactive plans, are (figure 3.6):

- $time_r$ - is the time needed by the agent between the momer contingency is detected, and until the proper reaction to it ca started; it depends on both the computational and non-computa resources of the agent, their capabilities and their load in that sit The value of this dimension grows with the number of t contingencies included in the reactive plan and with the complexi identifying them and their reactive responses.
- $resource_i$ - is the total requirement imposed on the agent's i-th re by the reactive plan containing the current contingency analyzed all the contingencies previously decided to be included for rea response and associated with this same situation. These dimension of special concern for real systems. Both computational a non-computational resources (including memory) are limited, and availability may be decisive for the successful completion of reaction (e.g., in the limit, a universal plan for a real domain require an infinite amount of memory, which is unacceptable in systems).

Inclusion of a reaction to a new contingency depends on the size c resulting reactive plan, which combines it with the set of all the reacti contingencies already decided to be included in the reactive plan fo situation. These contingencies were obtained from the agent's knowledg where they are indexed by their applicable situations, and have previously analyzed by this framework (since their criticality must be than the criticality of the currently analyzed contingency).

The agent's knowledge base includes all the contingency-reaction known to the agent, indexed by the situations in which they may appe with associated descriptions for the criticality space dimensions. We present in chapter 4 a formalism to construct languages for repres situations, contingencies and reactions in the knowledge base, designed advantage of the regularities of the application domain.

To continue with our example, the more contingencies (selected the 13 contingencies given in table 3.1) are included in the reactive pl more likely it is to decrease the responsiveness of the agent to each contingencies included. Since we have no information (yet) on the str of the reactive plan built by the reactive planner, and also on the resource limitations, we cannot actually specify how much each of the contingencies will increase the response time (we shall see in section that for some structures of reactive plans, adding some new contingenc in some circumstances, not increase the response time at all). In any wa agent will always try to include at least the reaction to child-in-front-of-the-car contingency, and will continue to add to it as as possible, in the order given in the table. However, it will not ad contingency if either (i) its estimated response time would be bigger th allowed response time, or (ii) if adding it would determine the response any previously included contingency to exceed its allowed response (given by the $Time_{rc}$ value of the criticality space associated with contingency).

Figure 3.6 summarizes the characteristics of the reactive plan s defined above, and their relationships (functions) to other elements framework. Functions f_t and all f_{ti} are explicitly contained in the reactiv planner model and are then used in conjunction with the limitations agent resources defined by the agent model.

3.2.5. Summary of the Framework

The purpose of our entire framework (and of the thesis for that m is to keep the reactive response time and other resources for very contingencies within acceptable (i.e. useful) bounds, while ensuring re behavior at least for the most critical contingencies known for situation. Given the information contained in the three spaces defined the agent has all the data it needs to be able, for every contingency, to decision of whether to include it or not in the reactive plan associated given situation. The result of processing the contingencies through the framework is a partition of the set of known contingencies possible in a situation into two classes: to be included in and to be excluded fror reactive plan.

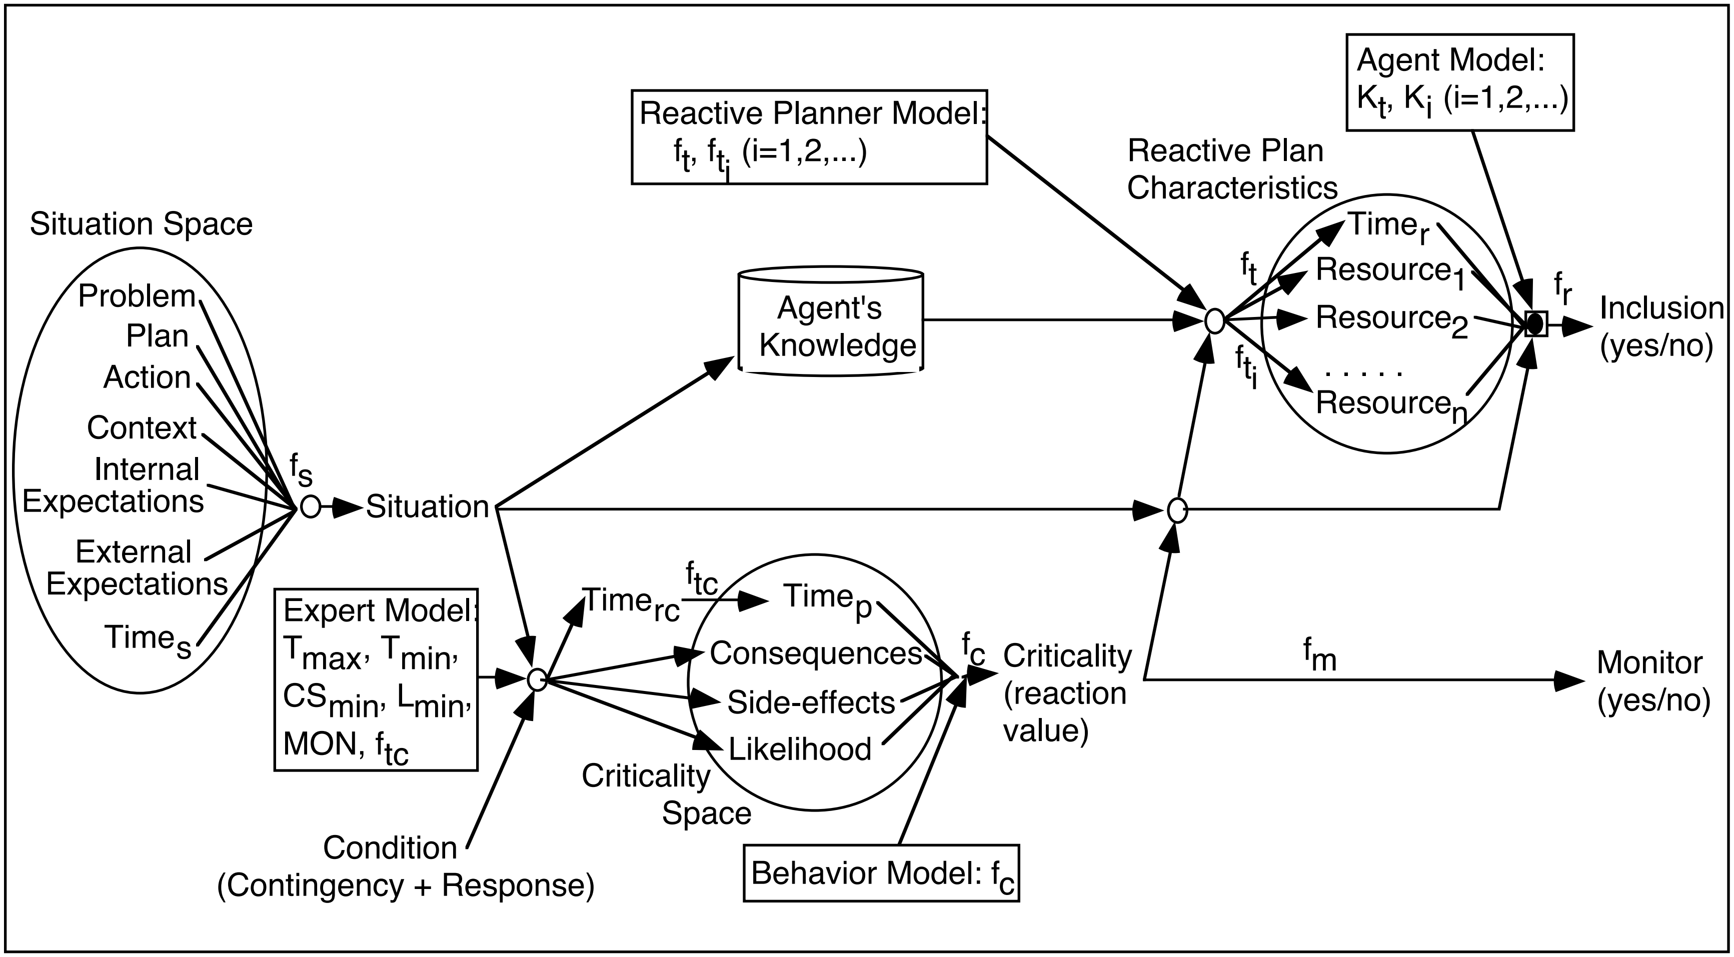

Figure 3.7. The Plan-to-React Decision Framework

Figure 3.7 shows a detailed summary of the framework for selectin contingencies for which reactions are prepared and those for w monitoring actions are added to the plan. It details the diagram prese figure 3.3, and essentially combines figures 3.4, 3.5 and 3.6. At any tim agent knows of a set of contingencies and reactions to them. Each conti may be associated with several regions in the situation space, and each in the situation space may have several contingencies associated (many relationship). Each contingency is characterized in a situation by a cri point. While the criticality value alone decides which contingencies wi monitored in which situations, the decision for including the treatment contingency in the reactive plan associated with that situation is made on both the criticality value, and the reaction value of the entire r plan for that situation, in relationship with the reactive planner mod the agent model.

Situation = f_s (Problem, Plan, Context, Action, Internal_expectations, External_expectations, Time_s)

Time_{rc} = f_1 (Situation, Condition)

Consequences = f_2 (Situation, Condition)

Side-effects = f_3 (Situation, Condition)

Likelihood = f_4 (Situation, Condition)

Time_p = f_{tc} (Time_{rc}) = k / Time_{rc}

Criticality = f_c (Time_p, Consequences, Side-effects, Likelihood)

Monitor = f_m (Criticality) - Expert Model

Time_r = f_t (Situation, Criticality, Agent's_knowledge, Reactive_planner_model)

Resource_i = f_{ti} (Situation, Criticality, Agent's_knowledge, Reactive_planner_model) (i = 1,2,...)

Inclusion = f_r (Time_r, Resource_1, . . . , Resource_n, Agent_model, Situation, Criticality)

Figure 3.8. Functional Relationships for the Plan-to-React Decision Framework

The set of functional relationships among the elements of framework is summarized in figure 3.8.

Appendix 1 presents the general agent architecture and the basic flow during the plan modification process.

Our agent integrates reactive responses with the plan to compensa the unfeasibility of universal plans. It does not only try to prepare f most frequent or likely contingencies, but also for some very infrequen which are very critical. Due to real-world resource limitations, some c frequent but not very critical contingencies may be excluded from reac favor of less frequent but very critical ones.

Space
Dimensions

Problem
Deliver package to work

Plan
drive car

Context
school time (May, week)

Situation
Action
drive straight, 25 mph

Intern. Expectations
reaching school

External Expectations
children in sight

Time
max. 3 mins.

Contingency

Child / Ball in front of car

Time
to avoid collision (short)

Consequence
fatal (very high) / small

Criticality
Side_effects
high

Likelihood
medium

React. Plan
Time
N.A. / to be considered

Characts.
Memory
N.A. / to be considered

Figure 3.9. Example for the driving domain

Two advantages of the framework introduced here are: (i) specification is general, domain and agent-independent, so we expect it applicable to a wide variety of agents working in a variety of environ and (ii) it is highly parameterized, which ensures a proper adjustment framework to a specific agent and to domain-dependent require (domain, expert, reactive planner, and agent characteristics and capab as well as to the desired type of behavior. In chapter 5 we claim and that the framework, as presented here, is free of redundancies; that is, the elements included in our framework are necessary to completely d the characteristics of a contingency and its reaction in order to allo agent to decide at planning time whether to prepare for the reaction contingency in that situation. While we cannot prove that the framew also sufficient (i.e. that there are no other elements needed for this d besides the ones described here), the experiments described in chapter successfully conducted using this framework. Should the need to exter framework arise, we believe that it can be easily done, while preservir elements and their structure discussed here.

Space
Dimensions

Problem
inguinal hernia

Plan
surgery procedure H

Context
heart disorder history

Situation
Action
apply anesthetic

Internal Expectations
get patient asleep

External Expectations
surgeon perf. incision

Time
from action to sleep

Contingency

heart failure

Time
to restore heart (short)

Consequence
fatal (very high)

Criticality
Side_effects
very low

Likelihood
high

React. Plan
Time
N.A. (irrelevant)

Characts.
Memory
N.A. (irrelevant)

Figure 3.10. Example for the anesthesia domain

Figure 3.9 presents a summary of the car driving example u throughout this section to illustrate our framework. Figure 3.10 prese example from a different domain - anesthesiology, to show the general our theoretical framework.

The agent is an anesthesiologist preparing for an operation du which contingencies that endanger a patient's life may appear. The si space is defined by the general characteristics of the operation (in hernia to be treated through a specific surgery procedure performed patient with heart disorder history). The plan analysis is at the point anesthetic is applied. This action will give rise to two kinds of expecta (milestones) to be watched for: as a result of the action, the patient sho asleep after a certain amount of time, and from the external environm expectation of an incision being performed by a surgeon. At this poin anesthesiologist agent analyzes as a possible contingency a heart failu has a short deadline (the time to restore the patient's heart without brain damage) and the consequences of not reacting in time are fatal high). It also has a high likelihood of occurrence, given the patient's m history. As we shall see in the following sections, since these characte yield a very high criticality value for this contingency, the agent probably decide to add monitoring actions to the plan, and will pr include its reaction in the reactive plan for this situation, almost regarc the rest of the contingencies relevant to the same situation (analogous child contingency in the driving example). In chapter 6 we present a set of results which we have obtained from our experiments in this m domain.

3.3. Establishing the Value of Reaction

As mentioned in the overview of the framework which we made section 3.2.1, our framework has two critical phases: establishing criticality (or reaction value) of the contingency, and making the decis whether to include its associated reaction into the reaction plan built current situation. In this section we will concentrate on the first of phases, and will leave the second one for the next section. But before w present our method for establishing the reaction value of a contingen have to talk briefly about the expert model, since it is according to s model that the values for the criticality space dimensions are specified.

3.3.1. The Expert Model

The situation-dependent criticality space values for contingency-reaction pair are supplied by an expert, and are thus sub the personal interpretation of the expert, according to his expert model. As our experiments have shown (chapter 6), the experts need not be precise in the absolute values they provide. It is enough if they are correct order and approximately of correct relative values. This is becau method for computing the criticality value (section 3.3.2) and the wa value is used further in the framework are robust (i.e. noise tolerant), the entire framework very robust. We shall substantiate these remar chapter 6, when we shall discuss the experiments we have conducted. these relaxed precision requirements, the experts with whom we have on the knowledge acquisition part of our experiments were able to s quickly and with very little effort suitable values for the characteristics contingencies in these experiments.

The values specified by the expert for each contingency are the time interval allowed between the moment a contingency is detected an its reaction is started, the consequences of not reacting to the contin the side-effects of executing the reaction associated with the contingenc the likelihood of occurrence of the contingency in that situation. Th three values are real numbers in the interval [0,10]. The values for th pressure dimension are positive reals; the upper limit for the time p depends on the threshold values imposed by the expert model, whic presented below. All these values may be specified qualitatively (e.g. fo consequences dimension using {very small, small, medium, high, very and are then translated into numeric values (e.g., corresponding t previous set of qualitative values, these numeric values will be in intervals: {(0.2], (2,4], (4,6], (6,8], (8,10]}. As seen in previous chapters, values are situation dependent; they may be different for the contingency associated with different points in the situation space.

The expert model reflects the expert's interpretation of the domain the way he or she estimates the values of the contingency character This model must include the following threshold values, which will be u the next section in our analysis:

- T_{max} - is an upper limit on the reasonable values for the time pr exerted by contingencies on the agent. A time pressure higher this value makes the reaction useless since it can only be taken too (the agent has no way to react before the deadline). In our dr example, the meteor contingency has a too short deadline to responded to realistically, so the agent is better off by not inclu such a reaction in the reactive plan (and leaving the reactive plan for contingencies that can be responded to in reasonable time).

- T_{min} - is a lower limit on the time pressure values for which the should try to respond reactively. If the agent has more time than threshold, then it can probably dynamically replan its response, leaving room in the reactive plan for other, more time pressu contingencies. Therefore, the value of reacting here is significa lower, although not zero - if the agent has left enough execu resources, then maybe it is still a good idea to prepare a rea response for such a contingency. For example, if the agent drivi car detects a traffic jam, it does not have to react (well, usually... can take its time to replan an alternate route. However, we can e imagine traffic jam situations in which the agent is much better o first reacting (and, say, leave the freeway) and then replanning, just by taking its time to dynamically replan (and, say, pass the fr exit).
- L_{min} - is a lower limit on the likelihood of occurrence of continge for which the agent should prepare reactions. A likelihood value than this threshold indicates that the contingency is so unlikel appear in this situation that the overhead of preparing and mana reactive response is probably unjustified, so the value of reacting is significantly lower. An example here can again be the met contingency, and maybe the airplane landing contingency too. treatment can be dangerous in certain domains where consequences may still be fatal, but in such cases this threshold ca lowered to zero. Also, the value of reacting if the likelihood drops l the threshold is again still positive (though much smaller), so if agent has left enough execution resources, then it may again be a idea to prepare a reactive response for such a contingency.
- CS_{min} - if the side-effects of a reaction to a contingency outweigh consequences of not reacting by more than this value, then i probably wiser not to take any action. In this case, like in the u time pressure threshold T_{max}, the value of reacting to the contingency is considered zero. An example is the contingency of a ball poppin in front of the agent's car: the side-effects of taking the recomm dangerous maneuver outweigh by far the consequences of hitting a at 25 mph, so the agent is better off by ignoring this contingency the reactive plan preparations.
- MON - is a criticality threshold beyond which monitoring actions fo contingency should be included in the main plan (even if reactions cannot be included); the reason is that the decision to include a re for a contingency is taken dependent on the agent's run-time res and performance, which may change over time, but are not taken account at this stage of the decision process. Also, these monito actions may detect a contingency for which no reactive response prepared, but for which the agent has the resources to dynam replan its response.

The agent model must also specify the function f_{tc} which transf real-time values into time-pressure values. These pairs of values are in proportional, so this function has the form:

$$Time_p = f_{tc}\ (Time_{rc}) = k\ /\ Time_{rc}$$

where only the constant k has to actually be specified by the expert moc has to be in some (weak) correlation with the two time pressure thre presented above.

Also implicitly contained in the expert's model are the functions f_1, which associate the values for the criticality space dimensions with eac condition-situation, as discussed in section 3.2.2.

3.3.2. Value of Reaction

The criticality value for a contingency-reaction pair is a measure c merit of the reaction to the contingency as opposed to dynamically rep a response to that contingency, in a particular situation in which contingency is known to possibly appear. This value induces an order r on the set of contingencies that can appear in that situation. This order to allow the selection of those contingencies that should be reacted to the limited resources of the agent. Function f_c, which computes the cri value for a contingency given the values of the characteristics of criticality space for the contingency, implements the evaluation functi the behavioral model to be exhibited by the agent.

The behavior model represents the type of behavior which the ager attempts to simulate. By imposing an order (i.e. a preference of treatme the set of contingencies associated with a situation, the agent commits it a pattern of reactive behavior. It involves both which contingencies preferred over which, and which contingencies are ruled out altogethe the reaction process. Each behavior model is characterized by an evaluation function which, given a set of conditions (pairs contingency-reaction) situation in which they apply, computes a score with the following pro the higher this score is, the better (more appropriate) that se contingencies is (according to the particular reaction philosophy of behavior model). The evaluation function orders the set of conting associated with a situation according to their priority for a reactive res

The behavior model is implemented in our framework through relative values of the parameters in the function computing the val reaction (which is presented here), and through the values of the thr on the criticality space dimensions (presented in the expert model) rela the values of the parameters of the criticality function. In chapter 5 we a few properties of the relationship between the evaluation function behavior model and the criticality function defined below. The most im property is that both functions define the same order relation on a contingencies associated with a same situation, which implies that criticality function can be consistently used to implement behavior moc

The criticality function we have used in our experiments has following general form:

Criticality = f_c (t, c, s, l) =

if (t > T_{max})
then $f_c = 0$
elseif (c + CS_{min} - s < 0)

then $f_c = 0$
elseif ($t < T_{min}$)
then $f_c = {}^{p4}*(c+CS_{min}-s)^{p5}*l^{p6})$
elseif ($l < L_{min}$)
then $f_c = {}^{p4}*(c+CS_{min}-s)^{p5}*l^{p6})$
else $f_c = t^{p1}*c^{p2}*s^{p3}*(c+s)^{p4}*(c+CS_{min}-s)^{p5}*l^{p6}$

where, for the purpose of stating the criticality function in a more su form, we made the following notations for the (situation dependent) cr space dimensions:

t = $Time_p$	(is the time pressure)
c = Consequences	(of not reacting)
s = Side-effects	(of the reaction)
l = Likelihood	(of encountering the contingency)

Parameters T_{max}, T_{min}, CS_{min}, L_{min} are dependent on the domain and are defined by the expert specifying the domain knowledge. Their m has already been defined in the previous subsection. They are import implementing a specific behavior model. For example, if the upper thr on the time pressure T_{max} is made lower, than more contingencies will be le out of the reactive plan since the agent estimates that there is not enou at execution time to give a timely response to these contingencies. behavior simulates the resignation behavior model [FAA, 1991] (the leaves responses to contingencies to others, since it believes there is no try to react to them, i.e. it believes that there is no time to take care o anyway). On the other hand, taking $T_{max} = \bullet$ emulates a behavior intended to avoid legal liabilities by always doing something.

Parameters p_1 to p_6 are also used to model different (human) behavio their relative values place the agent in different behavioral models and viewed as labels for human reactive behavior. For example $p_5 > p_6 > p_1 > p_2$ (with p_3 and p_4 very low) represents what is usually accepted as n behavior in the car driving domain: most importance is given to the pressure and then to the difference between consequences and likel with more emphasis on consequences; lastly, it also considers the likelih occurrence. Another behavior model in which consequences and esp side-effects are almost disregarded with respect to time pressure impl an attitude of invulnerability - the agent is prone to risk taking and d believe that anything wrong can happen to him. Again, it is importa notice the robustness of our model: the only important thing about parameters are their relative values, and these can themselves vary while still obtaining consistent results. This property makes the life c domain experts participating in the knowledge acquisition and behavior specification process much easier. In chapter 6 we shall discuss a num experiments we have made and how they justify our claims for the fran robustness.

As stated before, the value of reaction associated with a conting induces a total order relation on the set of contingencies associated v certain situation. This is only a partial order on the set of all conting known to the agent, since contingencies in different situations may (although sometimes can) be comparable according to their criticality This order relation is defined as:

"A is more_critical_than B " if and only if:

A and B are contingencies applicable in the same situation S, anc

A has higher criticality value than B, or

A and B have same criticality, but A has higher consequences,

A and B have same criticality and same consequences, but A has higher likelihood.

This ordering characterizes the behavior model of the agent. It subsequently be used to choose the contingencies for which reaction prepared (section 3.4.3).

Different combinations of these parameters defining the critic function are used in both the theoretical and experimental evaluati prove certain conjectures. In chapter 5 we claim that the paramet function defined here can implement the human reactive behavior described in the literature, and while we cannot formally prove this cla justify it through the experiments discussed in chapter 6. Therefore framework can also be used in psychological studies of "hazardous" atti certain high-risk domains like nuclear power plant operation and ai flying. In section 6.3 we present and briefly evaluate a series of exper we have conducted with our framework to simulate a number of re behavior models described in the literature.

3.4. The Reaction Decision Making

Making the actual decision of whether to include the contingency its associated reaction into the reaction plan built for the current situ the second and last critical phase of our framework. This phase is based the elements and the information previously acquired and computed framework. As shown in figure 3.7, there are two agent dependent mode participate in this phase: the reactive planner model and the agent They synthesize the agent's properties and the limitations on its resou planning time and execution time respectively. We first make a presentation of these models and the information they are expected to and then we give the actual algorithm for deciding whether to plan to r

3.4.1. The Reactive Planner Model

The reactive planner model describes the planning time properti the agent, and the characteristics of the reactive plans built by the age their relationships to the agent's execution time resources (computation as well as other non-computational resources). This model must allow agent, at planning time, to estimate the variations in execution time r requirements with respect to the growth of the reactive plan, namely w number of contingencies and reactions included in the reactive plan. accomplished by the functions f_t and f_{ti} in figure 3.11 which depic entire decision making process presented in this section.

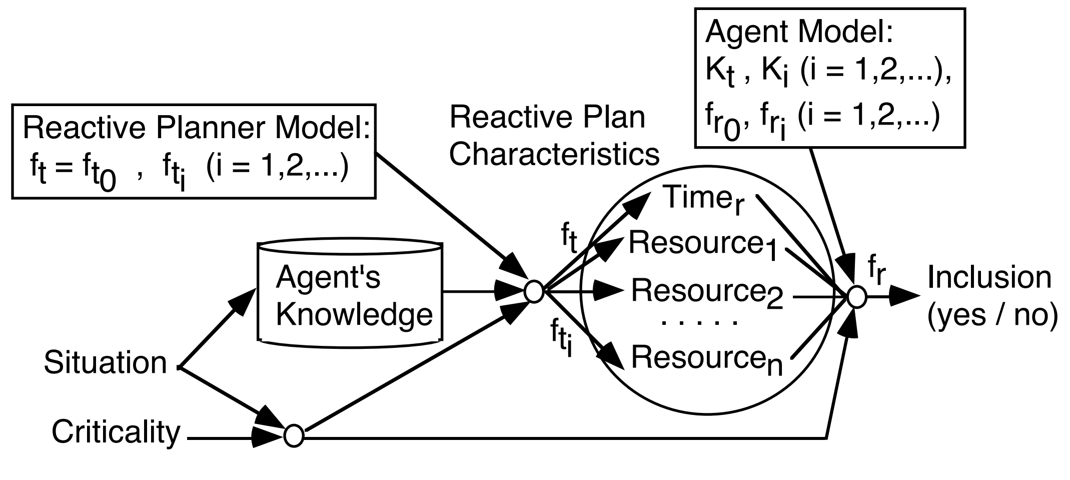

$\text{Time}_r = f_t$ (Situation, Criticality, Agent's_knowledge, Reactive_planner_model)

$\text{Resource}_i = f_{t_i}$ (Situation, Criticality, Agent's_knowledge, Reactive_planner_model)

Inclusion = f_r (Time_r, Resource_1,..., Resource_n, Agent_model, Situation, Criticality)

Figure 3.11. The Reaction Decision Making Phase

Function f_t estimates the time needed by the agent from the mom detects the existence of a contingency and until it can react to this pa contingency, when the reactive plan known to the agent in this sit contains the response to this contingency as well as responses to al contingencies with higher criticality which apply in the current situ The reactive planner model assumes that the agent can devote a computational resources to this task (this assumption is then taken care the agent model, described in the next section, which takes into accou overhead that the agent may experience in that situation). Functi estimates how much does the reactive response time increase, on avera adding this contingency to the reactive plan.

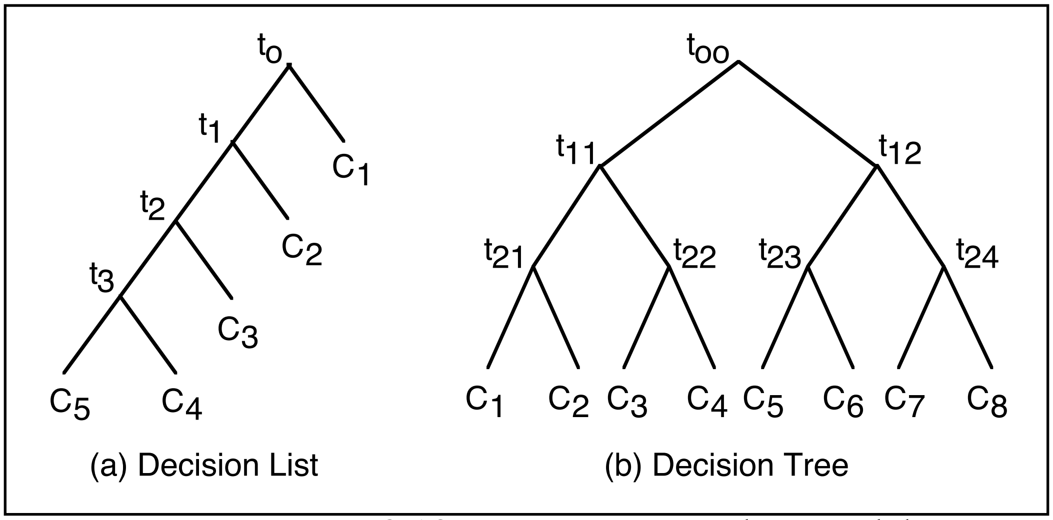

Figure 3.12. Two reactive plan models

Two commonly encountered examples of reactive planner models decision lists and decision trees. For a reactive planner based on decisio (figure 3.12.a), the time to react increases approximately linear wit number of contingencies to be considered, since for each new contin added to the reactive plan, a new test must be added to discrimina Therefore, the time needed to react to a contingency according to this will be the sum of the times required for each test that has to be done deciding on the contingency. If we assume the testing time to be ro constant, then the estimated time to react becomes:

$$Time_r = test\_time * rank\_in\_reactive\_plan$$

i.e. is directly proportional to the number of tests to be performed wl equal to the number of levels in the decision list before the continge question. In figure 3.12, t_i (i = 0,...,3) and t_{ij} (i = 0,1,2; j = 0,...,4) are test performed in order to determine the proper reaction to the contingen C_i (i = 1,...,8) are the contingencies (and their associated reactions) for the reactive plan contains responses.

If the reactive planner uses decision trees to index the reactions i final reactive plan, then the time to reach a response is closer to the lo of the number of contingencies (the base of the logarithm is equal t branching factor (assumed constant) of the decision tree), assuming ag approximately constant testing time. Figure 3.12.b presents such a co binary tree, for which the reaction time for each of the contingenc roughly:

$$Time_r = test\_time * \log_2 (number\_of\_contingencies\_in\_reactive\_pl$$

i.e. is directly proportional to the logarithm of the number of conting treated by that reactive plan (we assume complete decision trees, in wh k leaves (contingency-reaction pairs) are all situated at level m, or if $k = 2^m$, 2p of the leaves are at level 2^m and the other k-2p leaves are placed at leve $m-1$ when $k = 2^{m-1} + p$, $(1 < p < 2^{m-1})$.

Similar reactive planner models can be built for other method organizing the reactions in reactive plans.

Functions f_i have the same mission for each of the other critic resources of the agent (e.g. the amount of memory needed by the r plan, as well as any other non-computational limited resources that the might need in order to start its reactive response), as f_t has for compu time.

The two formalisms for structuring reactive plans mentioned a (complete binary decision trees and decision lists) deserve here a comparison. At the first glance, a qualitative reasoning seems to impl decision trees are better (or at least never worse) than decision lists. running the experiments described in chapter 6, we have found out that not necessarily the case. We shall show here when this is not necessaril and analyze and justify it. (A formalism is considered better if it can i more reactions to more critical contingencies in the reactive plan executed by the same agent with the same resource characteristics limitations, in identical situations). During this discussion we will assum all the tests require the same amount of time (T), and that there are tests available such that any arrangement of reactions in the resp reactive models is possible. In this case, responding to the n-th continge the reactive plan will take time $T * n$ in the decision lists case, and $T * \log_2 (n)$ in the case of complete binary decision trees.

We must note two things here: (i) different contingencies may significantly different time pressures (i.e. significantly different al response times), and (ii) a structural difference between decision list decision trees is that the complete decision tree takes the same amount to respond to all the contingencies, while decision lists respond fast contingencies placed towards the root of the list, and this response increases with the distance of the condition from the root.

Therefore, once the decision tree reactive planner has decide include a given contingency (say C) in the reactive plan, it can only a many contingencies to the plan until the estimated response tim contingency C becomes larger than its allowed response time. This mear the decision tree formalism is actually limited by the contingency wi highest time pressure which the agent decided to include in the reactiv This is not the case however for reactive planners based on decision Here, the contingencies with the highest time pressure can be placed t the root of the tree, and the response time to them will not be affected number of contingencies covered by that reactive plan. There contingencies with lower time pressure can still be added towards the the decision list, since they allow for a longer time of response, and wi affect the response time for contingencies placed higher on the list. A of experimental results which support this analysis (actually, as we earlier, they have prompted this analysis) are presented and discus section 6.2.

In summary, when the response times allowed by the continge under consideration vary within a small relative range, the decisior based reactive planner will be able to include more such contingencies all its leaves are reached in roughly the same amount of time). On the hand, when the time pressures of the contingencies vary widely (which to be the case in real-world domains), decision lists are better suite including responses to a larger number of contingencies, since testing

for contingencies with shorter time of response allows timely reactio
more contingencies with lower time pressure. Naturally, the best so
would be an incomplete decision tree which combines the advantages o
formalisms.

In this thesis, we assume that the agent has enough planning res
and time to build the most comprehensive reactive plans which do not
its execution time resource limitations. However, this framework may al
applied when dynamically replanning courses of actions, and wher
limitations on the agent's planning resources needed to build such r
plans may become a factor to be considered. In such cases, the re
planner model may also be required to estimate the complexity of the r
plan structuring algorithm. This estimate can then be taken into acco
our framework, and may lead to the decision of reducing the set of cor
to be included into the reactive plan, in order to ensure that the time
to construct the reactive plan will not exceed the time allowed for this ta

3.4.2. The Agent Model

The second agent dependent model involved in this later stage o
framework in which the agent makes the actual decision of whether to
the contingency and its associated reaction into the reaction plan built
current situation is the agent model. It synthesizes the agent's propert
the limitations on its resources at execution time.

The agent model describes the (situation dependent) respo
capabilities of the agent (figure 3.11). The functions (f_{ri}) describe the
variation of the availability of resource i (i=0 for computational time)
the fact that the agent cannot devote its entire resource r exclusively to
responding to that contingency. For example, the computational load
agent slows its responsiveness by a factor K_t greater than 1, and can be
expressed by:

$$f_{r0}\ (time_r) = time_r * K_t\ ;$$

or if the agent can devote itself to solving this contingency only after
constant time K_a, then

$$f_{r0}\ (time_r) = time_r + K_a,$$

and so on.

The agent model also supplies the amount of each resource (K_1, K_2,...)
that may be allocated to reacting in the given situation, for
non-computational resources. Example of non-computational resources a
the anesthesiology domain, oxygen masks and ventilators. Such resourc
available in limited quantity, and also may only become available af
certain waiting period. The agent model does not have to specify su
upper limit on the availability of resources for computational time, sin
is already specified separately for each contingency through the re
time allowed to respond to it (the time pressure dimension of the cri
space values associated with the condition in the agent's knowledge base

The agent model is very important in domains whe
non-computational resources may not be available all the time, but m
obtained after some waiting period (as in medical domains like anesthe
intensive care monitoring, or in nuclear power plant operation).

By comparing the requirements of each of the agent's run t
resources, for the set of the previously included contingencies plu

current contingency under consideration, with the limitations on availability of that respective resource (given by the agent model non-computational resources and the agent's knowledge base for time agent can decide whether this contingency can be included in the re plan for the current situation or not. We shall analyze this decision pro detail in the next subsection.

3.4.3. Deciding Whether to Prepare to React

The final purpose of this entire framework is to decide, for e contingency-response pair associated with a given situation, wheth preplan the reaction to it or not. As shown in figure 3.11, this decision i by comparing the estimated execution time resource requirements fo agent to respond to all the contingencies already decided to be included reactive plan plus the contingency currently under consideration, wi allowed response times for each of these contingencies in that situation.

Given the criticality of the current contingency and the set of the contingencies known possible in the current situation, this decision p proceeds as follows: the framework computes the agent's execution resource requirements to respond to any of the contingencies as:

$Resource_i = f_{ti}$ (Situation, Criticality, Agent's_knowledge, RP_model)

for each $resource_i$ (i = 0,1,...) of the agent (for a unitary exposition we sometimes call the agent's computation time as $resource_0$; all other re of the agent (possibly including the amount of memory needed by the plan, as well as other domain dependent critical and limited resource ventilators in an intensive care unit, etc.) are numbered starting with functions f_{ti} are given by the reactive planner model, and estimate increase in $resource_i$ requirements by adding this new contingency-reacti pair to the reactive plan. For i = 0, $f_{t0} = f_t$ estimates how much does the response time (considered from the time a contingency is detected, and reaction to resolve it can be taken) increase, on average, by adding condition to the reactive plan. As discussed in subsections 3.4.1, f_t is approximately linear for decision lists and roughly logarithmic for de trees. Obviously, the better the reactive planner model is (i.e. the bette estimates are), the better use of the execution time resources of the age be ensured by the selected set of contingencies.

As we have mentioned in section 3.3.1, the decision to monitor f contingency is taken based only on the criticality value of the contin and independent of the reactive plan characteristics. The reason is th decision to include a reaction for a contingency is taken dependent o agent's run-time resources and performance, which may change over but are not taken into account for monitoring purposes. Also, monitoring actions may detect a contingency for which no reactive re was prepared, but for which the agent has the resources to dynam replan its response. The decision to monitor is taken as a threshold func the criticality of the contingency:

Monitor = f_m (Criticality) = (criticality ≥ MON) =

| if (criticality ≥ MON) then f_m = yes

else f_m = no .

where MON is the monitoring threshold defined by the expert in the model (section 3.3.1).

The final decision of preparing a reaction for the currently ana contingency is taken by the function f_r:

React = f_r (Time$_r$, Resource$_1$,...,Resource$_n$, Agent_model, Situation, Critic =

if	criticality < MON	then	f_r = no
elseif	f_{r0}(Time$_r$) > Time$_{rc}$	then	f_r = no
elseif	f_{r1}(Resource$_1$) > K_1	then	f_r = no
elseif	f_{r2}(Resource$_2$) > K_2	then	f_r = no
.....			
elseif	f_{rn}(Resource$_n$) > K_n	then	f_r = no
else			f_r = yes .

= (monitor Λ ≤ K_i))) ,

where resource$_0$ is the real computational time$_0$ = K and Time$_{rc}$ is the real response time allowed by the contingency for the response to be s without consequences (the time pressure dimension of the criticality values for this contingency).

The functions f_{ri} are given by the agent model, and describe th execution time overhead imposed by other processes which the agent attend to in the same time in which it must respond to the contin Equivalently, they describe the availability of resource$_i$ for this reactive plan. They may be therefore situation dependent, and can be described as s the agent model. A common expression for these functions is of the form

$$f_{ri}\ (\text{resource}_i) = \text{resource}_i * k_t + k_a ,$$

where k_t is the overhead due to the agent's load (or the portion of it wh be expressed as a delaying factor), and k_a is an initial delay or cost associated with the use of that resource (for example, a process which cannot start a certain lead time, or a resource which cannot be delivered to the before a waiting period has elapsed). All these parameters must be speci the agent model.

```
//  input: a situation
//  output: a list of reactions (symptoms-actions pairs) for that situation

cr-list <- extract from the agent's KB all contingency-reaction pairs matching situation;
    //  cr-list is the set of all the contingencies known to the agent to be possible in situation

for each contingency in cr-list do
    time-pressure <- ftc (timerc);                                    // expert model
    criticality <- fc (time-pressure, consequences, side-effects, likelihood);  // behavior model
    if criticality > MON
```

```
        then  monitor <- true
        else  monitor <- false
    if  not  monitor
        then  eliminate this contingency from  cr-list
enddo
cr-list <- order cr-list by criticality value, then by consequences, then by likelihood
include <- ()
    //  include is the set of all the contingencies to be included in the reactive plan
    //        associated with situation
for   each contingency in cr-list do
    time_r <- f_t (include + contingency, situation) ;
    resource_i <- f_ti (include + contingency, situation);
    inclusion <- f_r (time_r, resource_1, ..., resource_k, time_rc, k_1, ..., k_k)
        //  f_r  returns true  iff there are enough resources to respond reactively to all
        //    contingencies previously added to the list include and to the currently
        //    considered contingency.
if  inclusion
    then   add contingency to  include
enddo
return  the list  include.
```

Figure 3.13. Reaction decision making algorithm

```
function fr (timer, resource1, ..., resourcek, timerc, k1, ..., kk)

    //  returns true  iff there are enough resources to respond reactively to all contingencies
    //  previously added to the list include, and to the currently considered one.

if  fr0 (timer ) > timerc
    then  return  NO;

for  i = 1 to  number_of_agent_resources do
    if   fri ( resourcei ) >  ki
        then  return  NO;
enddo

for each contingency in include do
    if   fr0 (contingency.timer ) > contingency.timerc
        then  return  NO;

    for  i = 1 to  number_of_agent_resources do
        if   fri (contingency. resourcei ) >  ki
            then  return  NO;
    enddo
enddo

return  YES.
```

Figure 3.13. Reaction decision making algorithm (continued)

One final set of parameters specified by the agent model are execution time resource limitations of the agent (K_i, $i = 1, \ldots$, in the formula for f above). They do not include $Time_{rc}$ which is a characteristic of contingency and therefore is specified in the agent's knowledge base. what the decision function does is simply to check that:

(i) the contingency is critical enough to be at least monitored for,
(ii) the agent will have enough time at execution to respond to contingency in the context of the larger set of contingenc considered for reactive response in the same situation,
(iii) none of the execution time limitations of the agent resources (b computational time) may be exceeded when attempting to respor this contingency, considering the entire reactive plan containir (i.e. all the contingencies with higher criticality, already decided t included in this reactive plan), and
(iv) the agent's run time resources are still enough to respond prop all the contingencies previously included in the reactive plan, this new contingency is added to the reactive plan.

This decision process ensures that no reaction is included contingencies which are not monitored for, and that there is enough a of each resource in order to attempt a reaction for all the conting included in a reactive plan. For the computational time resource, this that the time needed to start a reaction to the contingency is less than time allowed before the action must be taken (otherwise the reaction b useless).

Figure 3.13 makes a brief summary of the algorithm for deciding, a plan execution situation, on the set of contingencies to be included reactive plan which will be associated with the conditional plan befoı actual execution starts. The actual decision function f_r is presented seɪ in the second part of the figure.

The fourth test mentioned above essentially repeats the second third tests (carried out by the functions f_{ri}, i = 0,1,...) for each of the contingencies already decided to be included in the reactive plan. It m done each time a new contingency is considered for addition to the r plan, because the addition of the contingency, while possible from the p view of the restrictions imposed by its characteristics, may increase resource requirements to respond to previously included contingenci may therefore exceed the restrictions imposed by their critic characteristics. For example, in the case of a reactive planner base decision trees, adding a new contingency may force the reactive plan add one more level of tests in the decision tree, and thus increase the r time to all the contingencies included in this reactive plan. This way, sc them may now exceed the real time allowed for reaction to be taken, an reactions may become useless in that situation. (Conform to the analy section 3.4.1, the time to react to all the contingencies contained in a ı plan with a complete decision tree structure is approximately constar proportional to the depth of the decision tree).

The decision function f_r is applied in turn to each conting considered for the current situation, in the order given by their cri values, as defined in section 3.3.2 (each time, it applies each of the fu f_{ri} , i = 0,1,..., to each of the contingencies already included in the re plan and to the current contingency, considering the reactive plan to this contingency plus all the contingencies previously decided to be ir in the reactive plan for this situation). This iterative process is continue either all the agent's execution time resources are estimated to be exhau no more contingencies are known to the agent to be possible in the situation.

This concludes the presentation of our framework for deciding wł to plan to react. Given a plan situation and a set of contingencies known agent to possibly appear in this situation, it decides for which of contingencies the agent may prepare reactive responses, considerin execution time limitations on the agent's resources. In the next two cl we present a knowledge representation formalism to help the agent tc with the considerable amount of knowledge related to this decision p and theoretical justifications for some properties of our decision fram Then, in chapter 6, we present the results of our experiments usin framework. But before doing all this, let us see how the ideas presented can be applied to a related problem: given a plan situation and a contingencies known to the agent to possibly appear in this situation, for which of these contingencies the agent should prepare complete bı in the main conditional plan.

3.5. Conditional Planning

We briefly discuss here how the framework presented so far deciding whether to prepare to react to a contingency can be modif answer the question of whether the agent should prepare in its plan conditional branch for a contingency. We first resume our discussic section 2.1 regarding possible classifications of contingencies, and the adapt the previous framework to this new task.

3.5.1. Contingencies Revisited

In section 2.1 we have identified there types of contingencies that appear during the execution of a plan. They are classified according action taken by the agent at planning time to prepare for their occurr execution time. These types of contingencies are:

(i) contingencies for which the planner builds complete conditional branches, from the contingency state to the goal state, in the main As an example, suppose that the agent has two alternative route driving to work in the morning, depending on the color of a part traffic light when the agent reaches it: the regular plan assumes color is green, and the alternate branch is conditioned on the being red. For a non-driving commuter, the plan may involve wa or taking a bus, depending on the weather, and so on.

(ii) contingencies for which the agent prepares reactive responses, combined into reactive plans by a reactive planner, and attache appropriate segments of the complete plan provided by the cond planner. An obvious example is the one we used before, involvir child running in front of the car.

(iii) contingencies ignored by the agent at planning time; their tre at execution time can fall under two subclasses:

(a) dynamic replanning, if the agent has enough resources at exe time to perform it. As example, suppose that the agent encount traffic jam on a seldomly traveled route, for which it did not b to prepare a conditional plan branch before execution.

(b) noop, that is take no action, either because the consequences o contingencies are not high enough to warrant an action, or be the agent simply does not have the resources to take an actio solve them (e.g. they have a too short response time allowed). extreme example may be the contingency involving the me falling on the car, which we have encountered in table 3.1.

The justification for this classification is mainly related to the lir resources that a real agent can use. For a few contingencies, the ager generate complete plans and combine them in a conditional plan. Ho the agent's limited planning and execution resources do not allow for to contingencies to be treated this way. Still, the agent can prepare at pl time reactive responses for a larger set of contingencies; these response not ensure full solutions to the goal state, but they will give the age possibility to dynamically replan its actions at execution time. But in n can a real agent with limited resources prepare for all possible contin in a real world application domain. Many of these contingencies mu ignored at planning time.

Let us intuitively analyze now the characteristics of the exam given, and try to feel the qualitative differences among these class contingencies.

In the previous conditional planning example, the contingencies often, i.e. with a high likelihood (the occurrence probability may ap 50%, but should not exceed it, since if it does, then the contingency rather be considered the normal case and the main plan should be accordingly). Also, a solution to the contingency requires the preparat an entire plan branch all the way to the initial goal (since the executio may be critical and thus replanning cannot be used at any stage reaching the goal, i.e. a local situation stabilizing response to the conti is not sufficient), as well as certain resources whose availability mus planned in advance (e.g. an umbrella, or the correct set of maps fc alternate route to be traveled).

For the reacting case we have already devised a comprehen: framework stating the main necessary characteristics for a contingency considered appropriate for a reactive response. For the previous exampl include critical response time and high consequences of not respondir important characteristic is also that a short response (already availal sufficient to stabilize the situation and allow for replanning of the a actions all the way to the initial goal.

The rest of the contingencies will be ignored at planning time, bu have been able to further subclassify them. The ones for which the age: try to replan at execution time should not occur too often (otherw conditional branch may be appropriate), and should also allow for e time for the agent to be able to build the new course of action. Finall contingencies for which the agent will take no action anyway (e.g., falling meteor case) do not allow for enough time to respond to them, circumstances, given the agent's limited resources and execution capab:

In section 3.2.3 we introduced a criticality space, which is one po: representation of the space of contingencies, whose dimensions appropriate for reaction decision purposes. To facilitate the understan(the relationships among the classes of contingencies, we shall attempt simpler and more general graphical representation of the space contingencies, which can depict all the classes mentioned above. representation can conceptually be obtained from any more coı representation (like the criticality space mentioned before, or the imp space to be introduced later on in this section), by projecting the points space onto points in the simpler spaces defined here.

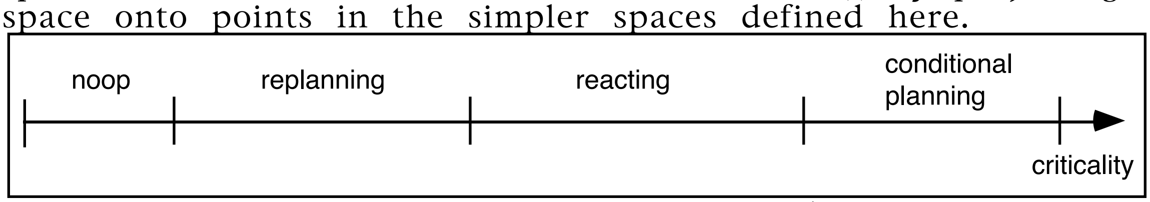

Figure 3.14. Contingency space - linear representation

The simplest representation for the space of contingencies is a l space in which contingencies are ordered by either criticality (as d before) or importance (as defined further in this section). Figure 3.14 that such a representation can outline the most frequent transitions l bordering classes, but cannot represent other still possible bordering between reacting and noop (e.g. determined by allowed response tim conditional planning and replanning (determined, for example, by planning time needed). Therefore, a planar representation (figure 3. more appropriate. The dimensions here are the reaction response valı the planning response value for the contingency. While much better representation still does not represent the direct relation between con planning and noop (which, to be fair, is the least frequent one, so representation can be used for most purposes). We have therefore dev third, 3-D surface representation using a spherical surface (figure 3.16 orthogonal dimensions (akin to latitude and longitude) are the same as second representation, and it can represent all the borders between p classes.

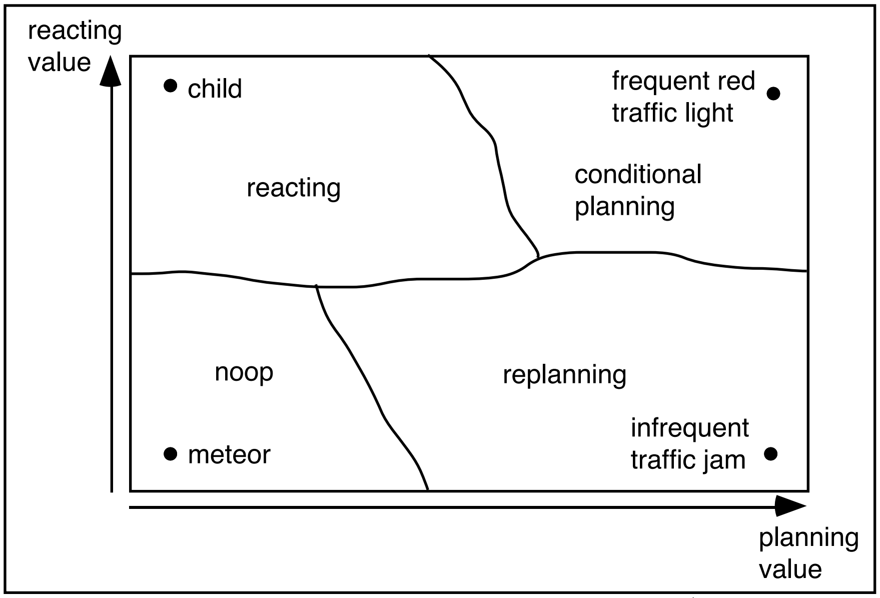

Figure 3.15. Contingency space - planar representation

The examples given with the informal description of these classes a beginning of this section constitute extreme cases in each class (figure In between these extreme cases there is an entire space of contingenci which more than one (in some cases even all) of the response altern may be justified. The borders among these classes in the space contingencies associated with a particular agent are determined by the resource capabilities and limitations. For example, conditional plannin replanning are separated mainly by the agent's planning resou replanning is circumscribed both by the agent's planning and exe capabilities, while reacting is mainly characterized by the agent's exe capabilities.

Due to the way the different classes of contingencies have been de in order to be able to best classify a given contingency, we only need membership decision frameworks for two of them, namely condit planning and reaction. We have already defined a framework for de whether the agent should prepare a reaction to a contingency in a situation. In the rest of this section we will give a description of a fram to decide whether to prepare a conditional plan branch for a contingen given situation.

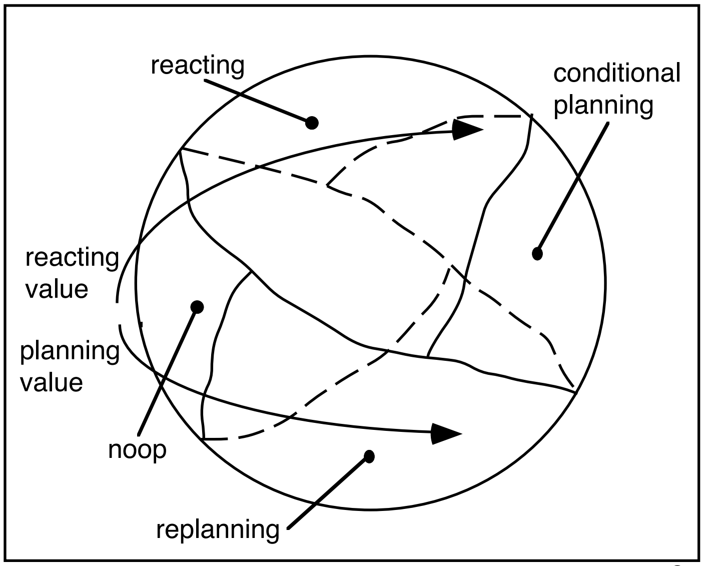

Figure 3.16. Contingency space - 3-D surface representation

There are two qualitative differences between conditional p branches and reactions. The first is that conditional plan branches re global solutions to the initial problem, that is, they are sequences of a which ensure that the agent reaches the goal (in the absence of contingencies). Reactions on the other hand are only single (or sequences of) actions, intended only to stabilize the situation so that th can then take its time to replan a solution from the state reached reacting to the initial goal. Therefore, on one hand reactions can be s the first steps of incomplete conditional branches, but in the same tim are more generally applicable than specific plan branches. There is al assurance that after executing a reaction, the agent may find a plan to g the initial goal, i.e. it is possible that the planner may subsequently fi solution from the state in which the agent finds itself after completin reaction to the goal; this is not the case for conditional plan bran assuming no other contingencies are encountered. Therefore, we a assume that a conditional planned branch is a better solution than a to the same contingency, and as a consequence, given a set of conting for a situation, the conditional planning decision framework should be before the reaction one.

The second difference involves the planning process itself. conditional planning, the planner has to work out a solution (sequen actions) from a given state (the contingency) to the goal. In rea planning, as assumed throughout this thesis, the agent already knows knowledge base) the best reactions associated with contingencies applicable classes of situations, so the only task of the reaction planner combine the reactions associated with the set of contingencies to be pr for, into a structure which will be conveniently searched at execution t determine the actual contingency encountered and its associated reactio decision trees, decision lists, etc.). Therefore, planning time is definite importance in conditional planning, but may not be an issue w structuring a reactive plan from a set of known reactions (if it cann ignored, then, as mentioned in section 3.4.1, the complexity of the re plan structuring algorithm can be taken into account in the Reactive I

Model, to further prune the set of contingencies for which reactions shc prepared).

Having noted these differences, we must now acknowledge that particular decision frameworks associated with the two classes contingencies have very similar underlying structures, so their preser may obey the same general lines. There are significant analogies betwe two problems and their solutions. They would suggest taking a ur approach and combine the two frameworks into a single one, with aest benefits of uniformity and elegance in presentation. However, we believ this would yield an unnecessarily complex framework, intuitively diffic present and understand. Therefore, as well as for easier understanding keep each framework manageable, we decided to present them separatel is also in agreement with the way in which an agent should apply although in different order. Indeed, the frameworks may indicate that contingencies are suitable for both conditional branch and rea preparation. In these cases a conditional branch should be prepared, sin assumed to be a more accurate solution, as argued before.

We first presented in sections 3.1 to 3.4 the reaction decision fram (which is the main topic of this thesis). In the remainder of this chap use analogies with the previous presentation to describe the condi planning decision framework, by pointing out their similarities differences. We transform one framework into the other by removing, and replacing some of its elements. Since the two frameworks are very cl form (although with underlying differences in content), an aesthe interested reader can easily merge them together if he or she so desires

3.5.2. Framework for Conditional Planning Decision

Let us first state the conditional planning decision problem, in a similar to the one used in section 2.2 for reaction. We assume the age built a linear main plan to go from an initial situation to a given goa issue then is to enable the agent, for each phase of the already built plan, to select the right set of contingencies for which to prepare conc branches all the way to the goal. That is, the **problem** is to specify a d framework which:

- given:
  - an intelligent agent with:
    - capabilities:
      - planning and dynamically replanning
      - monitoring
    - constraints:
      - limited resources
      - real-time performance
  - a linear plan by which the agent can achieve its goal
  - a set of contingencies known to possibly appear at certain ti during the plan execution, and for which the agent may pla conditional branch, each with:
    - known characteristics, associated with it (e.g. gravity consequences, time deadlines) and with preplanning conditional branch for it (e.g. resource requirements)
    - characteristics of their replanning alternatives (replanning and other resource requirements)

- enable the agent to decide for which contingencies to pre conditional branches in the plan (according to a desired beh pattern) while not exceeding the agent's planning capabilities preserving the real-time responsiveness of the agent to all t contingencies, given its limited resources.

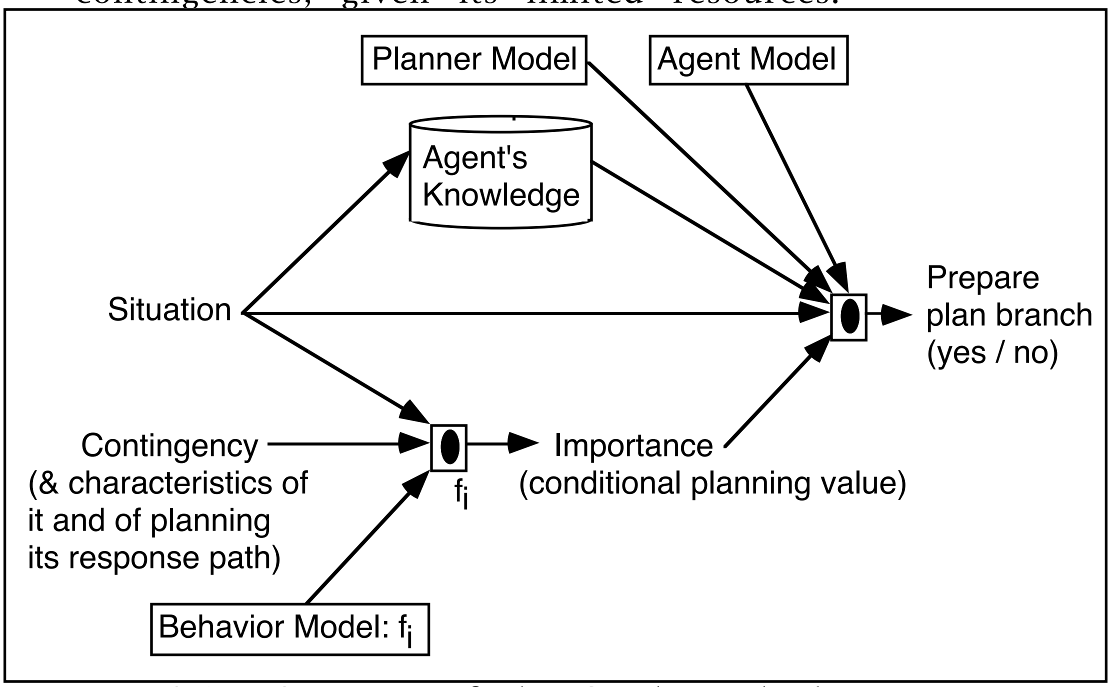

Figure 3.17. Overview of the Conditional Planning Decision Framework

As can easily be seen by comparing the two problems, they are si enough such that a solution to the second problem can be obtain relatively small modifications to the framework solving the first one. In the high level overview of the conditional planning framework show figure 3.17 is very similar in form to the one for the reaction fram depicted in figure 3.2. There are however a few underlying differences pointed out:

- the knowledge available to the agent and associated with contingency does not include the response to it, but only some g characteristics (outlined in section 3.5.3) of the planning process done for that contingency;
- the criticality (reaction value) computed by the reaction dec framework is replaced by an importance value (conditional pla value) which synthesizes how important it is for the agent to prep conditional branch for that contingency, i.e. what is the valu preparing a conditional branch for it in the plan vs. leaving it other possible treatments;
- the reactive planner model is replaced by a model of the conven planner used to build the initial plan and the conditional branche
- the final decision of the framework is now whether to prepare a b in the plan, instead of whether to include a reaction to the contin in the reactive plan associated with it.

Also, the agent model and the behavior model will reflect slig different characteristics in the two cases, and the functions used to c the conditional planning values and the final decision are based on so different variables, as will become evident soon.

Figure 3.18 presents in more detail the flow and source of inform through the new framework. Again, the comparison with the ge framework for the reaction case (figure 3.3) shows obvious simila between the two frameworks. The differences between the two framewo this level of detail and functionality are basically the same as the mentioned above for the higher level of abstraction used in the ove presentation.

Let us now briefly discuss each element of our new framework, compare it where appropriate to the equivalent element of the re decision framework. First, the situation spaces are identical in the frameworks, since a situation has the same definition and characte related to contingencies, regardless of the kind of response we prepa them.

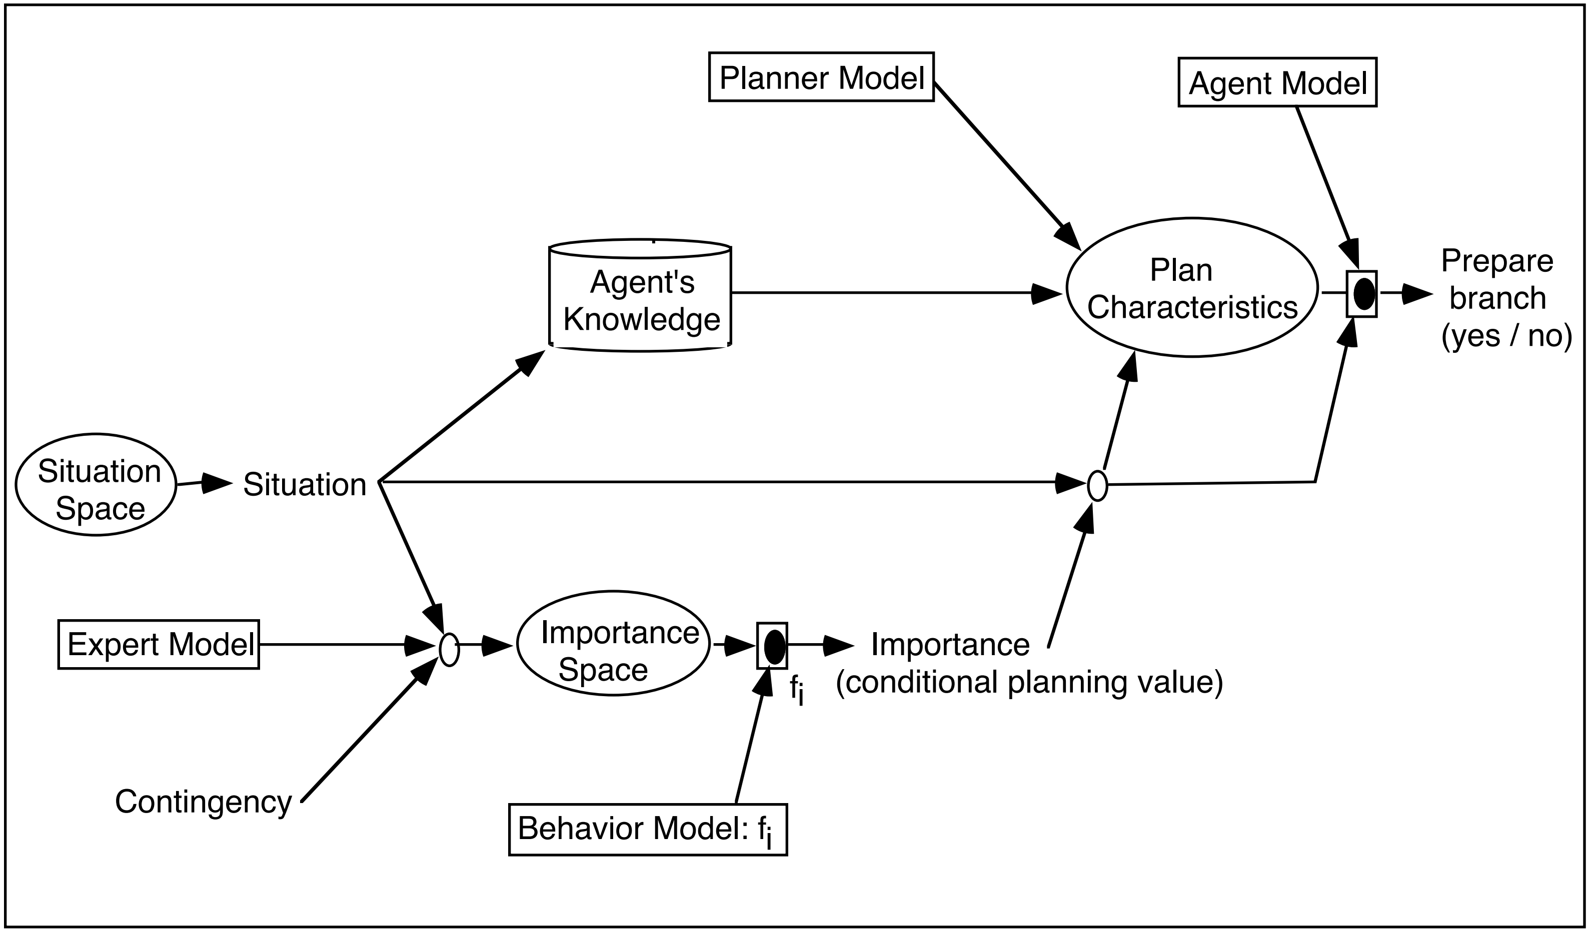

Figure 3.18. General Framework for Conditional Planning Decision

Two parts of the framework require special attention here. The establishes the conditional planning value of the contingency, and the takes the actual decision of whether to prepare a conditional branch f contingency. They are briefly discussed in the following two subsection then we conclude this presentation with a summary of the entire fra put together.

3.5.3. Establishing the Conditional Planning Value

Figure 3.19 presents the part of the framework concerned directly calculating a conditional planning value for the contingency in the situation. It is similar to figure 3.5 which shows the criticality space ar process of calculating the reaction value for a contingency. We concentrate here on the differences between the two frameworks at this

- the criticality space is replaced by an Importance Space which uses 5 dimensions to characterize a contingency from the conditio planning point of view. These dimensions are:
  - $Time_p$ - represents the same time pressure as in the reactive case; obtained from $Time_{rc}$ - the time allowed to respond to contingency, once an unexpected state is detected (same as in reactive case).
  - PTime - is the estimated planning time needed to build a branc this contingency at planning time (e.g., the time needed to pla alternative route, starting with a right turn at traffic light B, al way to the office); the simplest estimate may be, for example, planning time used to build the original plan from that point the goal.
  - Consequences - summarizes the consequences of not responding t contingency in the time allowed (same as in the reactive case).
  - PResources - is a measure of how hard (time consuming, ag resource consuming and any other costs involved) it is to obtai replanning time (during execution) the resources needed to r and carry out this plan branch (if not preplanned in adva Besides actual planning and replanning times, this also invo resources not needed in carrying out the initial plan, but which be needed for replanning purposes (like maps which may be ha obtain along the way) or for carrying out the alternate plan b (like an umbrella if it rains, or in medical domains a ventilato certain test results).
  - Likelihood - represents the likelihood of occurrence of t contingency in the given situation (same as for reaction).
- the Importance value which orders contingencies by their condit planning value (in the same way as criticality does for reaction).
- the function f_i() calculating the importance value for a contingency the form:

Importance = f_i (t, pt, c, pr, l) =

if (t > T_{pmax})
then $f_i = 0$
elseif (t < T_{pmin})
then $f_i =$

elseif ($pt > PT_{pmax}$)
then f_i =
elseif ($pr < PR_{pmin}$)
then f_i =
elseif ($l < L_{pmin}$)
then f_i =
else $f_i = t^{pp1} * pt^{pp2} * c^{pp3} * pr^{pp4} * l^{pp5}$

where, for the purpose of stating the importance function in a succinct form, we made the following notations for the (situa dependent) importance space dimensions:

$t = Time_p$, pt = PTime, c = Consequences, pr = PResources, l = Likelih

The two kinds of parameters involved are:

- (conditional) preplanning behavior model parameters: pp_1 to pp_5
- parameters specified by the expert model: T_{pmax}, T_{pmin}, PT_{pmax}, PR_{pmin}, L_{pmin}. They are domain dependent and are defined by expert specifying the domain knowledge. Their meaning is def below.

- the Expert Model reflects the new dimensions of the importance spac must specify the following:
  - functions:
    - f_{tc}: transforms (as for reaction decision) real-time values i time-pressure values, inversely proportional, so it has general form:

$$Time_p = f_{tc}\ (Time_{rc}) = k\ /\ Time_{rc}$$

  - parameters:
    - T_{pmax} - time pressure threshold - for greater time pressure, attempt of response is useless (akin to T_{max} for reacti decision);
    - T_{pmin} - time pressure threshold - for smaller time pressu dynamic replanning is possible (and thus less costly, since it be done only if the contingency actually arises); akin to T_{min} for the reaction framework;
    - PT_{max} - planning time threshold - if the estimated planning required is longer than this threshold, then the agent may no able to complete the conditional branch in the estima available planning time;
    - PR_{min} - replanning resources threshold - for smaller values, agent has enough execution time resources such that replan is possible (and presumably less costly);
    - L_{min} - likelihood threshold - if lower likelihood, the cost preparing a conditional branch for this contingency in situation is probably unjustified (same as for the react decision framework).
- the parameters of the Behavior Model (pp_1 to pp_5) also reflect the new dimensions of the importance space as well as the new func computing the importance value.

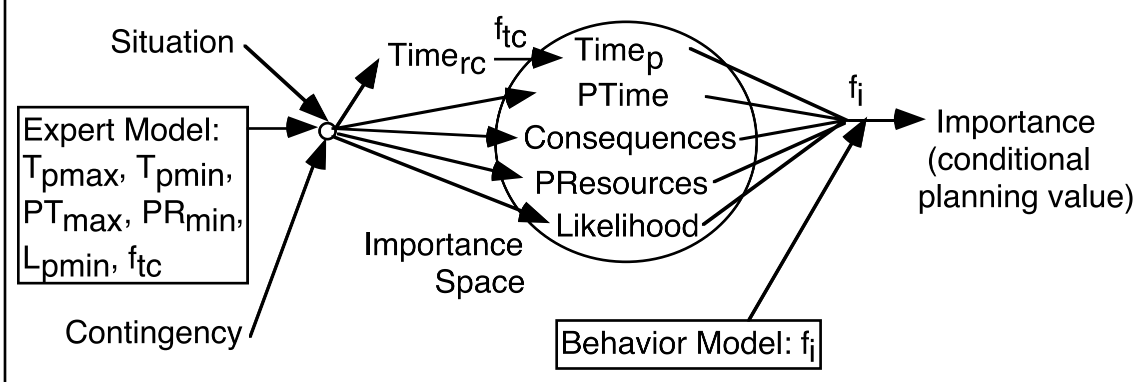

$Time_{rc} = f_1$ (Situation, Condition) Consequences $= f_2$ (Situation, Condition)
PTime $= f_5$ (Situation, Condition) Likelihood $= f_4$ (Situation, Condition)
PResources $= f_6$ (Situation, Condition) $Time_p = f_{tc}(Time_{rc}) = k / Time_{rc}$
Importance $= f_i$ ($Time_p$, PTime, Consequences, PResources, Likelihood)

Figure 3.19. Establishing the Conditional Planning Value

Note that the time to preplan a conditional branch may be diff from the time to replan it at execution time, because of different res availability and different information availability; in the driving exa when building the plan at home we may have all the necessary maps, s which may be unavailable when replanning later on during the execut the initial plan, and obtaining them may be time consuming, thus mak initial planning time shorter than replanning time. On the other hand replanning, the agent may have access to more accurate state infor than at initial planning time, and therefore the initial planning time this case be longer than the replanning time (for example, when the must replan its route due to a traffic jam, it has more knowledge about alternatives are available for faster traffic flow, than it could have bef actually reached this point in the plan execution).

Also note that side-effects are not taken into account in this frame since once prepared, the conditional branch is executed as a regula which under normal circumstances leads to the final goal (the side- were a measure of the risk of not being able to reach the final goal an once the reaction is executed).

3.5.4. Deciding Whether to Plan a Conditional Branch

Figure 3.20 presents the part of the framework concerned with the decision of whether to prepare a conditional branch for the continge the given situation. It is similar to figure 3.11 which shows the rea decision making phase of the previous framework. We shall outline he differences between the two frameworks at this stage:

- the reactive plan characteristics space is replaced by a Plan Characteristics Space whose dimensions characterize the en conditional plan to be built, from the point of view of the ag planning and execution resources. These dimensions are:
  - TPTime - measures the total planning time needed by the planner conditional branch for this contingency will be planned in add

to the main plan and conditional branches for the continge
already selected for conditional planning;

- $Time_r$ - is the estimated time needed by the agent to respon execution time, to this contingency, given that the conditional includes a branch for it together with branches for t contingencies already selected for conditional planning (simila the reaction framework);
- $Resource_i$ (i = 1,2,...) - represents the total requirements impose the agent's i-th resource by the conditional plan containin branch for this contingency as well as branches for t contingencies already selected for conditional planning (simila the reaction framework); an example of such a resource may memory amount required by the plan, which is separat represented in figure 3.20 by the total plan size (PSize).

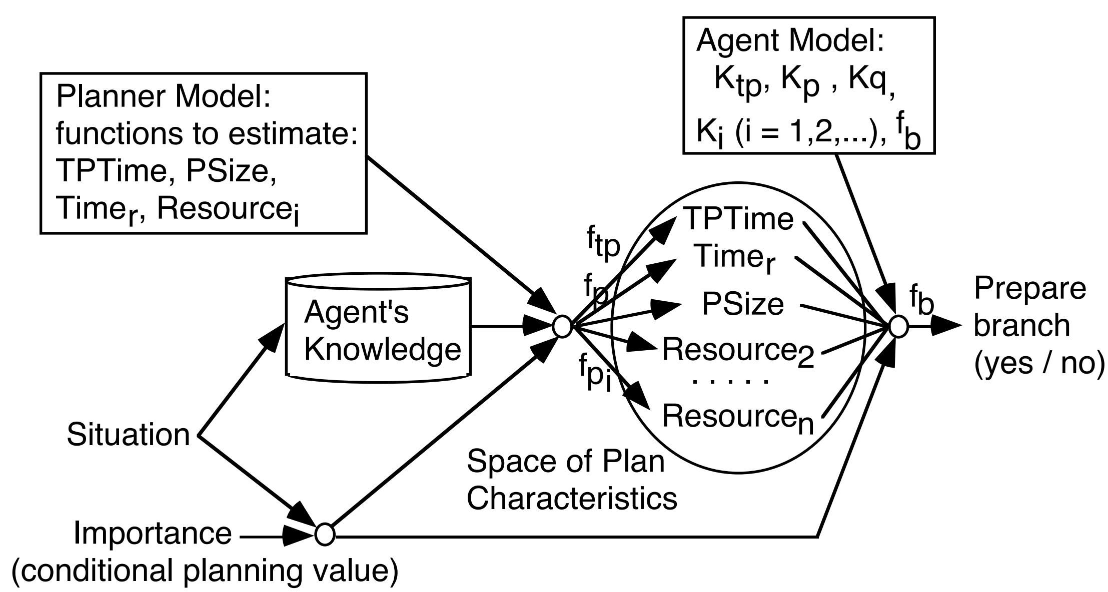

$TPTime = f_{tp}$ (Situation, Importance, Agent's_knowledge, Planner_model)

$Time_r = f_p$ (Situation, Importance, Agent's_knowledge, Planner_model)

$PSize = f_{p_1}$ (Situation, Importance, Agent's_knowledge, Planner_model)

$Resource_i = f_{p_i}$ (Situation, Importance, Agent's_knowledge, Planner_model)

$Prepare\_branch = f_b$ ($TPTime$, $Time_r$, $PSize$, $Resource_2$,..., $Resource_n$, Agent_model, Importance, Situation)

Figure 3.20. The Conditional Planning Decision Making Phase

- the Planner Model reflects the new dimensions of the pl characteristics space. It must supply the following functions to est values for these dimensions:
  - f_{tp} - estimates the time needed to build the plan, including a b for this contingency (in its simplest form, it may simply add already estimated times to build each individual branch);

- f_p - estimates the time needed to respond to the contingency whe plan includes conditional branches for it and for all continge with higher importance (similar to the reaction decisi framework);
- f_{pi} (i = 1,2,...) - estimates the amount of $resource_i$ needed to respo the contingency when the plan contains conditional branches f and for all contingencies with higher importance (similar to reaction decision framework); for i = 1, the function estimates amount of memory the agent needs in order to accommodate conditional plan.

- the Agent Model also reflects the new dimensions of the pl characteristics space. It must specify the following:
  - estimated maximum resource amounts that may be allocated by agent to this task:
    - K_{tp} - the maximum planning time allowed to build the conditi plan (i.e. before any execution begins)
    - K_1, K_2, ... - the maximum amount of $resource_i$ (i = 1,2,...) available at execution time (i = 1 for memory availability or, equivalen plan size)
  - functions to estimate resource utilization:
    - f_{bp} - the increase in planning time due to the agen computational overhead at the time of planning; it may be of form:

      f_{bp} (TPtime) = TPtime* K_p

      where K_p is a factor greater than 1, or:

      f_{bp} (TPtime) = TPtime + K_q

      if the agent can devote itself to planning for this continge only after some constant time K_q, and so on.
    - f_{bi} (i = 0,1,...) - the variation of the availability, at execution t of $resource_i$ (i=0 for computational time; i = 1 for memory or size) due to the fact that the agent cannot devote its en $resource_i$ exclusively to responding to that contingency (same the functions f_{ri} for the reactive plan characteristics space the reaction decision framework).
  - the function (f_b) making the actual decision for a conditional b preparation:

    Preplan = f_b (TPTime, $Time_r$, PSize, $Resource_2$, ... ,$Resource_n$, Agent_model, Importance, Situation) =

if	f_{bp}(TPTime) > K_{tp}	then	f_r = no
elseif	f_{b0}($Time_r$) > $Time_{rc}$	then	f_r = no
elseif	f_{b1}(PSize) > K_1	then	f_r = no
elseif	f_{b2}($Resource_2$) > K_2	then	f_r = no
.....			
elseif	f_{bn}($Resource_n$) > K_n	then	f_r = no

else f_r = yes .

= ≤ K_i)) ,

where resource$_p$ is the planning time, resource$_0$ is the execution real computational time, and $0 = K$ time$_{rc}$ is the real response time allowed by the contingency for the response to be started wit consequences (the time pressure dimension of the importance values for this contingency).

Figure 3.21 shows a detailed summary of the framework for sele the contingencies for which complete conditional branches are to prepared. We shall not continue the discussion on this topic, since this t mainly concerned with developing the reaction decision framework, a have included the presentation of the conditional planning framework point out that, after we have one of the two frameworks well defined experimentally proved adequate, the other one can be developed us certain degree of analogy.

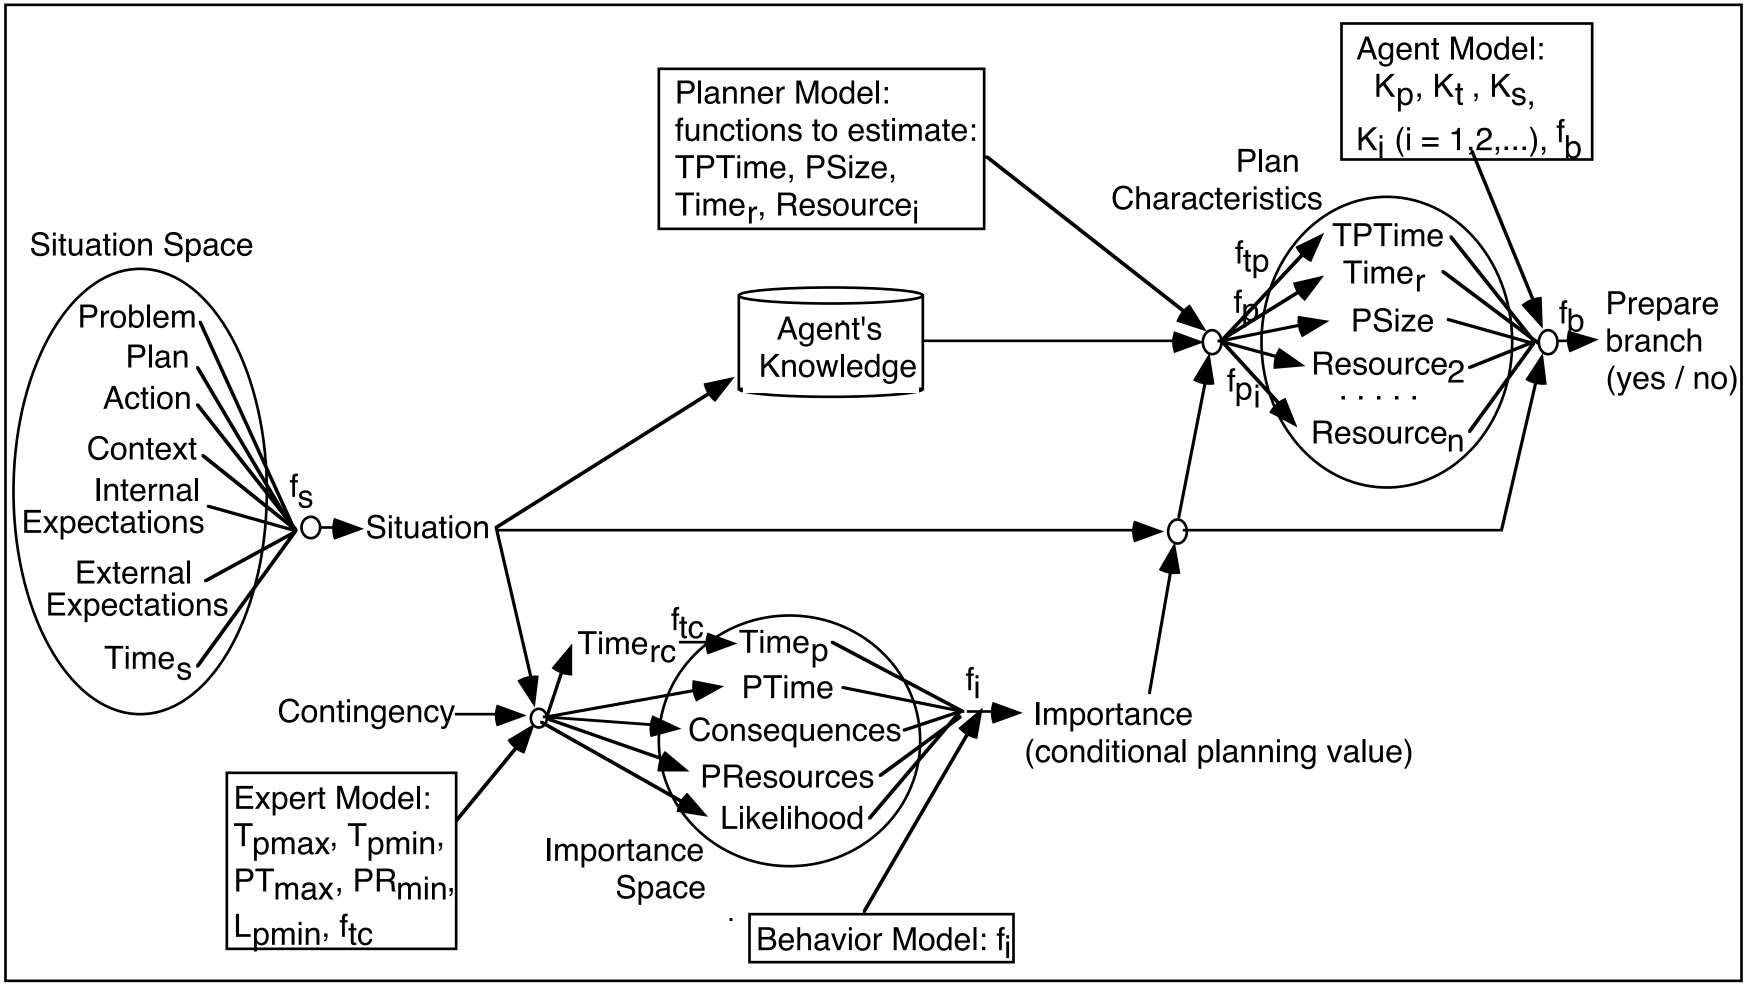

Figure 3.21. The Conditional Planning Decision Framework

Figure 3.22 presents two examples of contingencies that may wa conditional planning of branches to solve them. They are both taken fr driving domain, but may appear in significantly different circumstance they both largely illustrate the way the framework is intended to be ap

Space
Dimensions
Car driving to work
Car driving to Reno

Problem
Go from home to work
Go from Palo Alto to Reno

Plan
Drive car
Drive car on I80

Situ-
Context
Morning, commute time
Winter, night time

ation
Action
Approach intersection B
Approach Sacramento

Int Expect
Observe traffic light
See Sacramento

Ext Expect
Heavy traffic
Dark (night time)

Time
max. 3 mins.
30 mins.

Contingency

Red traffic light (slow - all following lights red too)

Cold & raining hard - maybe snow in mountains

Time$_p$
To reach intersection B
To reach junction I80, I50

Impor-
PTime
High (≈1/2 of main plan)
High (≈1/2 of main plan)

tance
Consequence
Late for imp. meeting
big delay, maybe life threat

PResources
Need maps + planning
Need maps + planning

Likelihood
High (≤ 50% of time)
High

Conditional
plan branch

Right turn at traffic light, then alternate route
Use I50 - longer but more reliable when snowing

Figure 3.22. Conditional planning examples

The first example is the one we mentioned in this section before: o usual commute to work, there is a certain traffic light which, if red on means that all the following traffic lights will be red, and the commut take significantly longer than if an alternate route is followed by mak right turn. However, this alternate route is slower if the traffic ligl question is found on green.

The second example is set during a trip from the San Francisco Bay to Reno at night time during winter. If it is cold and raining ar Sacramento, then there is a good chance that the usual (and faster) f may be closed in the mountains due to snow, so an alternate route is wi: it has to be prepared in advance since it may require maps for plannin;

A comparison between figures 3.21 and 3.7 shows that the frameworks are close enough so that an aesthetically concerned read(easily merge them into a single framework, so we shall not concern ou with this topic anymore. Instead, in the next chapter, we present a knc representation formalism to help the agent to cope with the consid amount of knowledge related to these decision processes.

Chapter 4
Knowledge Representation Formalism

In order to operate in an environment, the agent has to possess a knowledge about that environment. For the purpose of deciding whet plan to react to possible contingencies according to the framework pr in the previous chapter, the agent has to possess three types of infor knowledge about situations that may be encountered during plan exe knowledge about the contingencies that may happen in these situatior knowledge about the most suitable reactions to these contingencies. agent's knowledge base contains associations of contingencies and appropriate reactions. Each pair contingency-reaction is indexed ir knowledge base by the characteristics of the situation in which contingency may appear and in which that is the most suitable reactior Therefore, each condition stored in the knowledge base has three parts:

(i) a description of the contingency (signs, preconditions, and so on) set of values for the dimensions of the criticality space

(ii) a description of the best suited reaction for this contingency i situation described by the third part

(iii) a description of the situation in which this contingency may a and in which the best response to it is the reaction described in (ii). This description contains the values for each of the se dimensions of the situation space mentioned in chapter 3.

In the previous chapter we have presented the kind of inform associated with each of these classes of knowledge. With the exception contingency information which contains numerical values for the valu the characteristics of the criticality space, the rest of the informati symbolic. This includes the values for the situation space dimensions descriptions of the contingencies, and the descriptions of the actions make up the reactions to contingencies. Theoretically, one could us natural language to specify these values. However, such a natural lar interface and the mechanisms to process the information in such a for are beyond the scope of this work. In order to contain the explosi complexity generated by such a natural language representation, we defined a knowledge representation formalism which restricts the desc language for each of the classes of knowledge under consideration, retaining enough flexibility to be suitable to any domain and with the advantage of a well defined structure which can be used in the rea process.

In this chapter we shall discuss this knowledge representat formalism for each of the classes of knowledge involved, with examples the driving domain. We shall first present the general idea which is app all the three classes, and then we shall discuss an example of representi contingency description knowledge for the car driving domain. Apper presents an example of representations of reactions and representati situations for the same domain.

4.1. Description Languages

The need to devise a knowledge representation formalism for desc situations, contingencies and reactions has arisen from two consideratio

(i) the space of all possible natural language descriptions for these of knowledge is too large to be manageable; this in turn gener problems like the possibility of having different representations the same piece of knowledge and the associated difficulty of comp such representations and deciding on their identity. For example, car driving set of reactions we have used during the previous ch "steer" may be equivalent to "change direction", and clearly situation has many different equivalent ways of being described.

(ii) the practical application domains for the framework of dec whether to prepare to react presented before have a significant a of inherent structure implicitly contained in them and it woul unfortunate not to be able to exploit this structure. Notice for ex that eleven out of the thirteen examples of contingencies we gave the car driving domain (table 3.1) use the action "brake" in description of their associated reactions. The car driving domain also a significant amount of inherent structure in the description possible contingencies. For example, the following two contingen "Child runs from right, 20 m in front of car" and "Adult crosses street from right 20 m in front of car" have both the same criti space values, and the same associated reactions, and therefore d need separate representations in the agent's knowledge base.

If the structure of the application domain is not taken into accour explosion of the information that has to be recorded in the agent's kn base quickly exceeds any realistically manageable amount for a operating in the real-world domains described in chapter 2. For example are any number of individual situations for which the same contingency-response applies, and it would be entirely unreasonab represent each of them and all their associations with different conditi

Given these considerations, we have designed a representat formalism for these classes of knowledge which preserves most of flexibility of the natural language representation, while allowing the to take advantage of the structure of the domain.

For each domain there are nine languages which must be define language for describing the contingencies, one for describing the rea and seven languages for describing the values associated with each c seven situation space dimensions. Each of these languages will be des according to the same formalism, so we shall only describe the formalism and then (in the following section) we will give an example of each language in the driving domain.

The expert is required to define a hierarchical vocabulary for ea these languages in his domain. The words in the vocabulary are part into two classes: terminals and nonterminals. Each nonterminal repres class of words (both nonterminals and terminals). The terminals are clas themselves. The expert must also define all the membership and su relationships among the words in the vocabulary. Each such relatic defines a directed edge in a tree (actually a forest, since there is no ne full connectionism in a vocabulary) which induces a hierarchy ont vocabulary. The tree is actually an AND/OR graph, in which the OR r represent the membership and subclass relationships, and the AND represent structural relationships among words in a valid sentence language. Our formalism has a few common features with the lang representation formalism presented in [Utgoff, 1988], although it diff many other aspects. Our formalism defines a context free grammar:

$$G = (N, T, P, S), \quad \text{where}$$

- N - is the set of nonterminal words of the domain dependent voca defined by the expert
- T - is the set of terminals in the vocabulary
- P - is the set of productions of the grammar; there are two type productions:
  - unit productions, defined by a membership relation betwee terminal and a nonterminal or by a subclass relation among nonterminals
  - non-unit productions, defined by AND nodes in the vocabulary - give the rules of correct derivations in the language.
- S - the start symbol, which is either the root of the tree (if one exist if the vocabulary is organized as a forest, then it is a new nonter (OR node) to which all the roots of the trees making up the fores connected through subset relationships edges.

This context free grammar defines the language used for descr either the contingencies in the domain, or the reactions, or one of the characteristics of the situation space, with one important difference fr classic theory of context free languages: every word in the vocabulary n part of a sentential form in the language, that is, both terminals nonterminals may be used to build sentential forms. The set of terminals vocabulary makes up the agent language, that is, the set of all indi describable contingencies (or reactions, or characteristics of situatior sentential form containing only terminals represents a description specific contingency, reaction, or situation characteristic. It can als interpreted as a description of a singleton set of contingencies, reacti situation characteristics. A sentential form containing at least a nonte symbol represents a description of a set (of any cardinality) of contingencies, reactions or situation characteristics. This extension c context free grammar paradigm enables us to represent the structure application domain.

Our formalism also extends the classical context free grammar par with the notion of identification functions for nonterminals in vocabulary. An identification function is a compact way of represent large set of class membership relationships or a large set of sub relationships. For example, the nonterminal "slow_driving_speed" ca identified by a function defined as:

$$f\ (\text{speed}) = 5 \text{ mph} < \text{speed} < 20 \text{ mph}.$$

This function replaces all the edges in the tree between the nonte "slow_driving_speed" and all the terminals "speed = x" where x can t the discrete values representable in the machine (or in the defined vocabulary) between 5 mph and 20 mph.

Every tree generated by a vocabulary as described above defines partial order relations among the words of the vocabulary as well as the set of sentential forms that can be built. The elementary partial relation, which we call "contains", among words in the vocabulary, is define as: "a contains b" if and only if a and b are words in the vocabulary, a is a nonterminal, and either a and b are identical, or a contains b as a memb is a terminal), or a includes b as a subset (if b is a nonterminal). The extended partial order relation with the same name is applied to sentential through the following definition: "A contains B" if and only if A and B are sentential forms in the language (according to the previous definition every word in A contains the word in B in the corresponding position, i.e

if $A = a_1 a_2 \ldots a_k$ and $B = b_1 b_2 \ldots b_n$

then a_1 contains $b_1 b_2 \ldots b_{p1}$, a_2 contains $b_{p1+1} b_{p1+2} \ldots b_{p2}$, and so on until a_k which contains $b_{pk-1+1} b_{pk-1+2} \ldots b_n$.

In the next section we shall give an example of applying the form described here to the car driving domain and we shall present the voc trees which can be used to express the contingencies given in table 3.1 effectiveness of this representation formalism in structuring the appl domain will be illustrated by the realization that the same vocabular allows for the representation of a much larger set of contingencies, essentially the same knowledge acquisition effort and similar storage computational requirements. We shall then conclude this chapter with summary of the advantages of this knowledge representation formalism

4.2. Example

In this section we shall present the hierarchical vocabulary consequently the grammar) which are sufficient to represent the tl contingencies for the car driving domain listed in table 3.1. Appen contains a description of the vocabulary for representing the reactio those for the situations encountered in chapter 3. The vocabularies w only be able represent the knowledge contained in table 3.1, but also more.

Figure 4.1 presents the hierarchical vocabulary for represen contingencies.

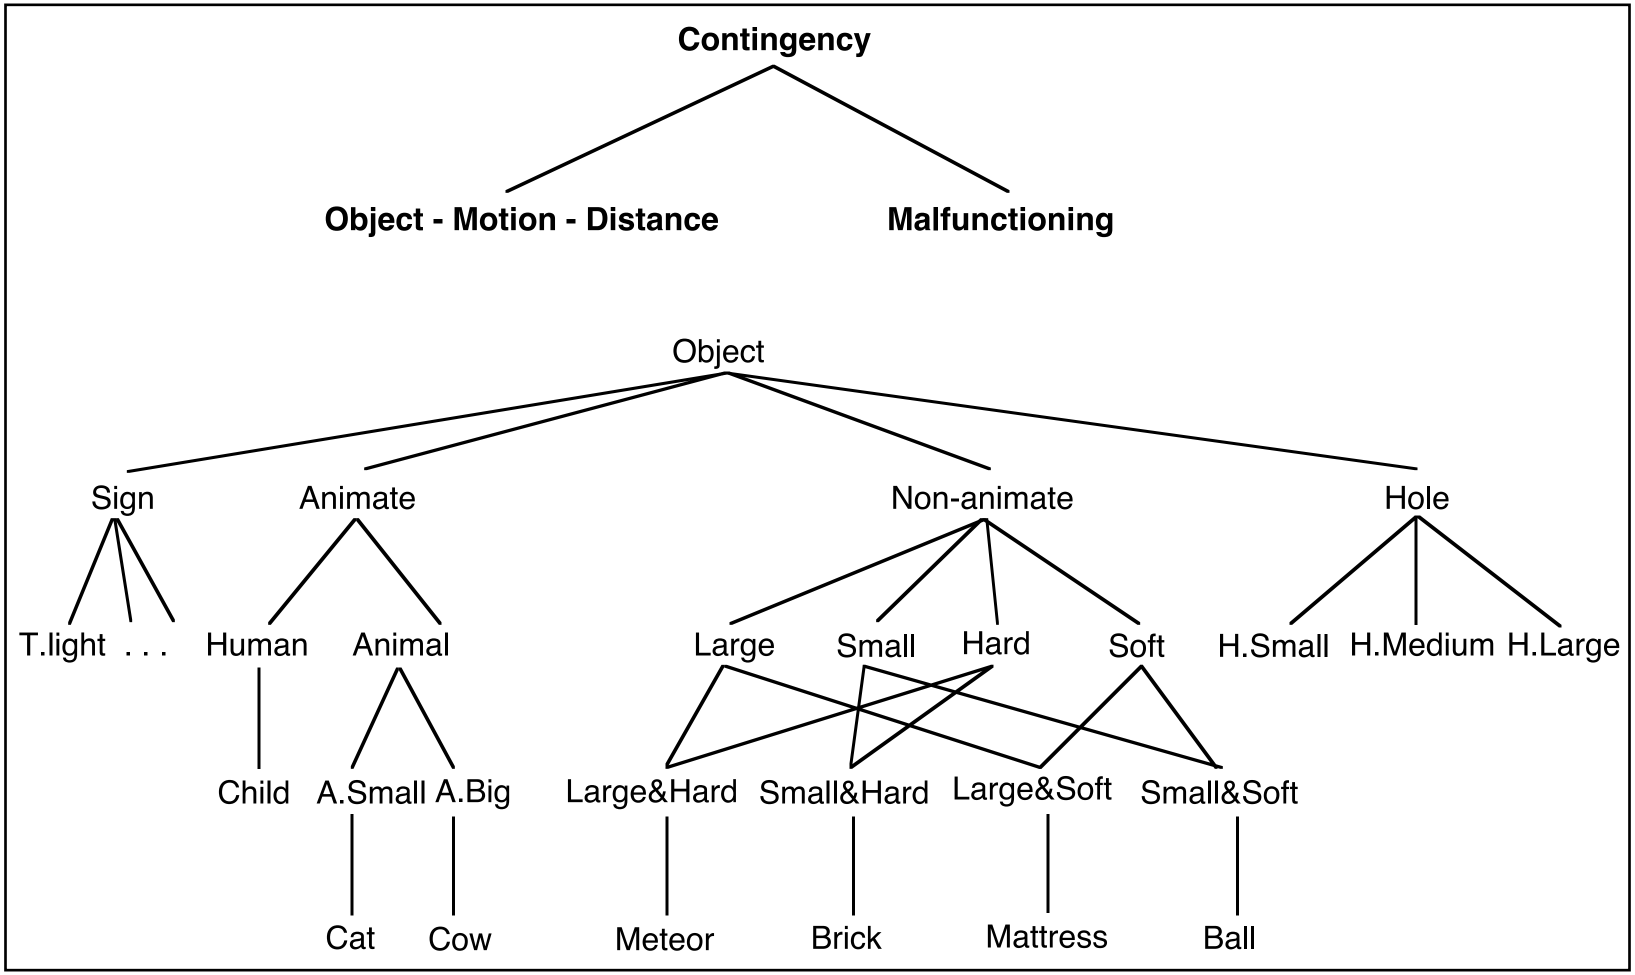

Figure 4.1. Vocabulary for describing contingencies in the driving domain

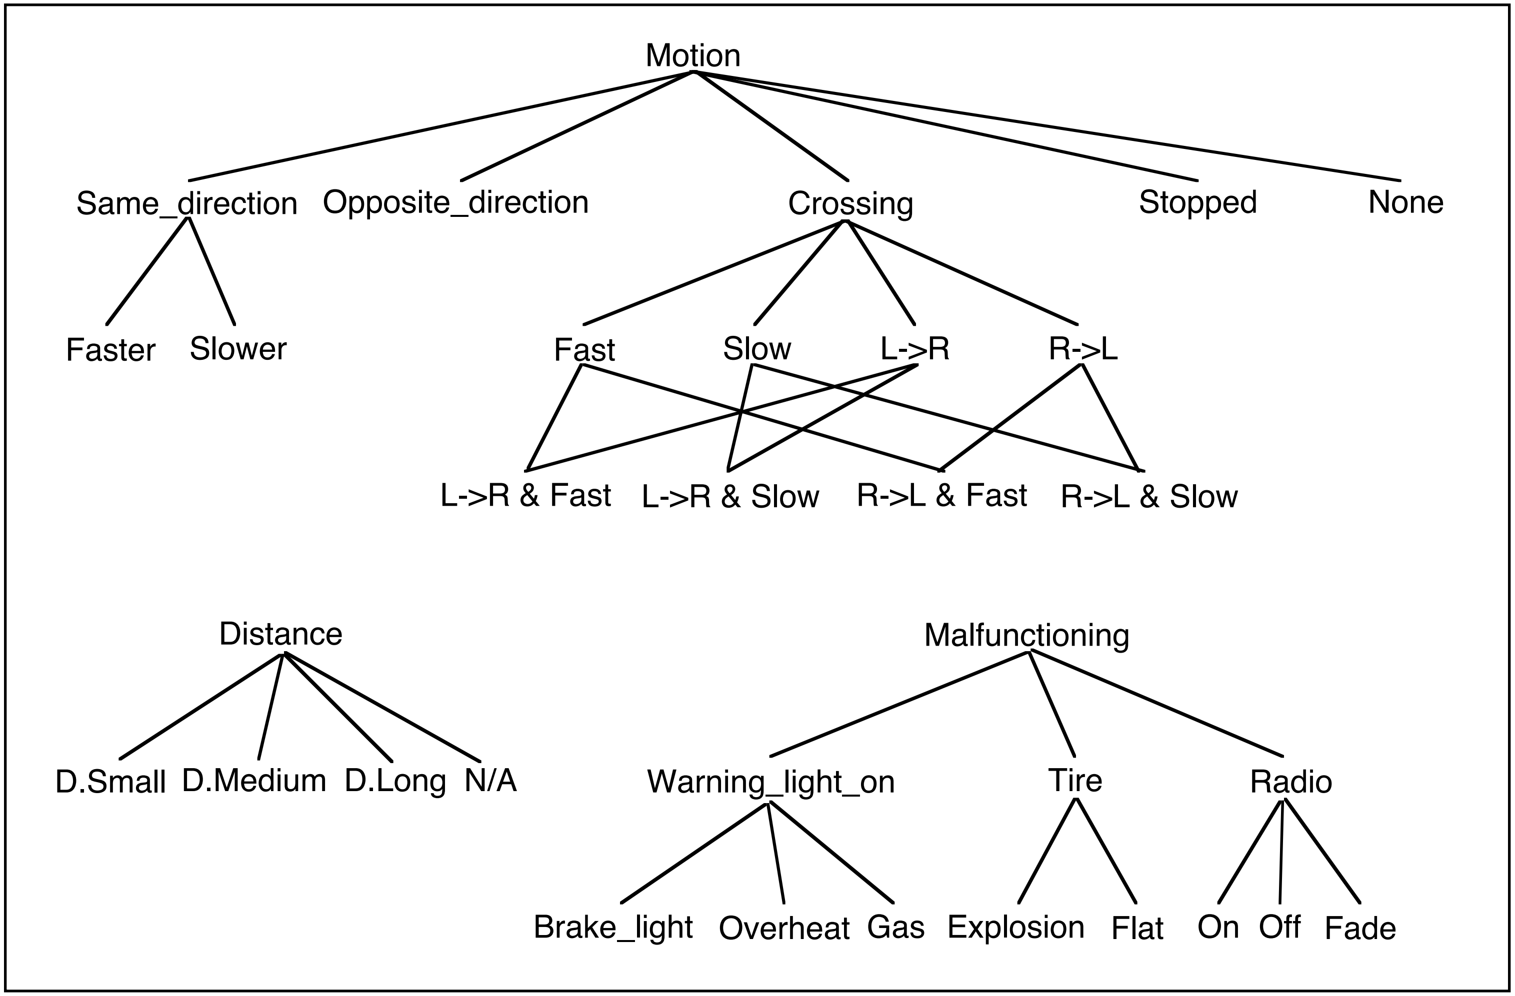

Figure 4.1. Vocabulary for describing contingencies in the driving domain (continued)

This hierarchy is equivalent to the following grammar:

G = (N, T, P, S), where:

N = { Contingency, Object, Motion, Distance, Malfunctioning, Sign, Animate, Non-animate, Hole, Human, Animal, Large, Small, Hard Soft, A.Small, A.Big, Large&Hard, Small&hard, Large&Soft, Small&Soft, Same_direction, Crossing, Fast, Slow, L->R, R->L, Warning_light_on, Tire, Radio }

T = { T.light, Child, Cat, Cow, Meteor, Brick, Mattress, Ball, H.Small, H.Medium, H.Large, Faster, Slower, Opposite_direction, L->R&Fast, L->R&Slow, R->L&Fast, R->L&Slow, Stopped, D.Small, D.Medium, D.Long, Brake, Overheat, Gas, Explosion, Flat, On, Off, Fade }

P = { Contingency -> Object - Motion - Distance | Malfunctioning
Object -> Sign | Animate | Non-animate | Hole
Sign -> T.light | . . .
Animate -> Human | Animal
Non- -> Large | Small | Hard | Soft
Hole -> H.Small | H.Medium | H.Large
Human -> Child | ...
Animal -> A.Small | A.Big
Large -> Large&Hard | Large&Soft
Small -> Small&Har | Small&Soft
Hard -> Large&Hard | Small&Hard
Soft -> Large&Soft | Small&Soft
A.Small -> Cat | . . .
A.Big -> Cow | . . .
Large&Hard -> Meteor | . . .
Small&Hard -> Brick | . . .
Large&Soft -> Mattress | . . .
Small&Soft -> Ball | . . .
Motion -> Same_direction | Opposite_direction | Crossing | Stopp
Same_direction -> Faster | Slower
Crossing -> Fast | Slow | L->R | R->L
Fast -> L->R&Fast | L->R&Slow | R->L&Fast | R->L&Slow
Distance -> D.Small | D.Medium | D.Long | N/A
Malfunctioning ->Warning_light_on | Tire | Radio
Warning_light_on -> Brake_light | Overheat | Gas
Tire -> Explosion | Flat
Radio -> On | Off | Fade }

S = Contingency

Some derivations may be done through identification functions. example, the grammar symbols D.Small, D.Medium, D.Long can be consi nonterminals (instead of terminals like in the previous example), an actual values of the distance can be considered terminals. Then, a fu like:

D.Small = 5 m < distance < 25 m

can be used to perform the transition over the edge linking D.Small wi actual terminal, say "distance = 20 m".

Every contingency in table 3.1 can now be obtained through a nu of different derivations in this grammar, and since the reactions to usually apply to more general contingencies, the derivation can be stop

higher levels, since a sentential form can contain both terminals
non-terminals in the grammar. For example, the contingency:

"Child runs from right 20m in front of car"

can be obtained through the following derivation:

Contingency ->
Object - Motion - Distance ->
Animate - Motion - Distance ->
Animate - Crossing - Distance ->
Animate - Crossing - D.Small ->
Human - Crossing - D.Small ->
Human - Fast - D.Small ->
Human - R->L&Fast - D.Small ->
Child - R->L&Fast - D.Small ->
Child - R->L&Fast - distance=20m.

Any sentential form encountered during this derivation (or durin other derivation leading to the same contingency) can be used to denc contingency. Each such sentential form contains (and denotes) the set contingencies derivable from it. The same reaction specified for contingency in table 3.1. (Brake hard and steer right) would probat recommended for the entire set of contingencies: "Human - R->L&Fa D.Small", while the consequences of the contingency would probably ha same value for an even larger set of contingencies: "Human - Crossi D.Small".

Clearly, this small vocabulary is not enough to describe all pos contingencies in the driving domain. It was not our goal to provide s vocabulary and grammar. However, while every contingency in table 3. be derived in this formalism, it supports the derivation of many contingencies for the driving domain. In fact, just by enlarging the terminals, the number of contingencies expressible with this small gr becomes very large indeed. This fact underlines the most important ad of this representation formalism, namely imposing a (hierarchical) st on the set of possible contingencies in the domain, which then makes much easier to be stored, managed, analyzed and reasoned about.

The knowledge representation formalism used in this chapter allow collapsing entire sets of contingencies in categories, thus alleviating problem of knowledge base size explosion.

Another advantage of this representation formalism is that it ca used in a future work for learning purposes, that is for learning which contingencies are similar from certain points of view of the ger framework for deciding whether to plan to react introduced in this which contingencies have the same characteristics, or the same reactio may appear in the same situations. Concept learning mechanisms ([Mi 1978; Mitchell & al., 1983; Dabija, 1990]) can be applied to continge represented in this formalism, mainly because the terms "classification and "concept description" used in machine learning are synonyms wit description", which represents any sentential form derivable in formalism. This representation can also be used to discover new class contingencies (non-terminals in the vocabulary) which have eludec expert's attention when specifying the domain, through bias shifting automatically [Utgoff, 1988], or interactively with the expert [Dabija & 1992a,b]).

The primary disadvantage of the knowledge representation form described in this chapter is that the expert must define the structure domain, that is, must specify both the nonterminals of the grammar (r the terminals), and the membership and subclass relations among elements of the vocabulary. This may place some burden on the expe may make the knowledge acquisition process more difficult. Anc disadvantage is that each specified vocabulary is domain-dependent (an user-dependent), as are all the relationships expressed through formalism. They all reflect how the expert who participated in the kno acquisition process views the domain. But the advantages (mentioned ab structuring the domain and significantly reducing the size of the kno base outweigh by far this disadvantage, with the added benefit that the is himself compelled to structure his own knowledge of the domain. problems may further be alleviated by using the learning techn mentioned above: some of them will attempt an automatic restructuring knowledge base, while others will help the expert to structure his knowledge of the domain through interactions with the system. Howev knowledge acquisition work has been done as part of this thesis.

The entire previous discussion applies equally well to represe reactions and situations. Hierarchical vocabularies may be used to cl reactions since in real domains there are usually a small set of actions can be combined to produce useful reactive plans, which are then ass with classes of (rather than individual) contingencies. This allows a structuring of the set of reactions, which in turn ensures better analy facilitates the reasoning about different sets of related reactions and characteristics with respect to the framework presented in the pre chapter.

The same is true for representing situations. Here this represent formalism is even more useful since the variety of situations in real d in virtually infinite, so any mechanism which induces a certain structu facilitates the reasoning process is more than welcome. Identific functions are also particularly useful here, since the values of some d dimensions may belong to continuous sets. Classes of situations de through this knowledge representation formalism and satisfying "contains" relation, are used to more efficiently index contingencies reactions in the knowledge base (as opposed to indexing them to sp situations, which would be prohibitive in any reasonably-sized real do The vocabulary for representing situations may be partitioned into distinct vocabularies, one for each dimension of the situation sp Alternatively, for uniformity of presentation reasons, we can combin seven vocabularies into a single one, with a new start symbol, "Situation" by adding to the grammar a new production of the form:

Situation -> Problem - Plan - Context - Action - Internal_Expectations
External_Expectations - Time,

where Problem, Plan, Context, Action, Internal_Expectations, Time a External_Expectations, were the start symbols for each of the vocabular the seven dimensions of the situation space.

The hierarchical vocabularies (and the grammars they generate representing the reactions listed in table 3.1 for the car driving doma for representing certain situations in this domain (including the one throughout chapter 3) are presented in appendix 2. Some derivatio

sentential forms encountered in chapter 3 for reactions and situations
driving domain are also discussed in appendix 2. As in the case
contingencies, these vocabularies can represent a much larger se
reactions and situations than the ones we have encountered durin
presentation in the previous chapters, with very little or no overhead
once again supports our claim regarding the power of the know
representation formalism presented here, and outweighs by far
disadvantages.

Chapter 5
Theoretical Analysis

The dream of any designer is to prove that his product is the ideal solve the original problem that motivated the design. In our case, this mean proving that the framework introduced in chapter 3 is always a decide, for any given situation, which of a set of contingencies possible situation should be selected at planning time to prepare reactive respon It would also mean to prove that this is the simplest framework with property, and also that the set of contingencies selected by it make th possible use of the agent's execution time resources. But since our objec to design a framework that is applicable in the most demanding re domains, theoretically proving all the previous properties is beyond means. However, we have been able to theoretically justify some of properties and some weaker versions of others. For the rest, while w believe that they hold in our case, we could only provide experim justifications which are presented in the following chapter.

In this chapter we present the theoretical justifications for a few c properties stated above. We first justify (through counterexamples) our that each of the elements included in the framework is necessary, that the framework is free of redundancies. Next we claim that the framewo consistently implement desired behavior models, and that the crit function defined in section 3.3.2 can implement any known type of r behavior; we formally justify the first of these claims, and in the next we present an experimental justification for the second one. Finally, w claim that the set of contingencies selected through our framework mak optimal use of the agent's execution time resources while simulating desired reactive behavior pattern, and we formally justify it. One more which cannot be justified theoretically but is verified experimentally next chapter is that the knowledge required by our framework in or execute properly exists and can be acquired in real domains.

But in the next section let us first briefly review the gene assumptions of our framework, which will be used during this chapter. following sections we shall then present our theoretical justifications t properties of necessity, consistency and optimality of the framework.

5.1. Assumptions

As discussed in chapter 2, during our presentation we have r certain assumptions about the problem we attempted to solve. T assumptions refer to both the agent, and the environment in which designed to work. The assumptions regarding the agent refer both t agent's execution capabilities, as well as to the design of its different modes.

The main assumptions for designing our framework were:

- about the agent capabilities:
  - planning (and planning to react)
  - monitoring
  - reacting
  - limited resources (including computational time)
- about the task environment:
  - real-time requirements
  - complex - there exist a large (infinite) number of possible situatic
  - complex - there exist a large (maybe infinite) number of poss contingencies in each situation
- about the agent control modes:
  - planning is better than reaction, whenever the resources (incl execution time) allow it
  - planning (like reaction) is useless whenever there is insufficient to reach a solution
  - reaction is faster than planning
  - limited resources allow only for limited amounts of reaction

We also assume that the agent's knowledge base always conta whatever information may be necessary for the operation of the fram Whether such information exists in real life and whether its acquisiti the knowledge engineer or the agent is possible will not be of concern chapter. However, we claim that this information indeed exists and acquisition is not very difficult, and we support our claim with experiments described in the next chapter and performed in different requiring quite different types of human expertise.

Note that all the assumptions listed here are not very restrictive. In they mostly restate the applicability conditions for our framework, pr in chapter 2. This means that the following results do not lose their ge from these assumptions.

Any other local assumptions that we shall make in order to allow perform theoretical analyses of our particular claims will be stated wh they apply.

5.2. Necessity

We claim that each element of our framework is indispensable for final decision, that is that each element in the framework is necessary final decision, or alternatively, that the framework is free of redund The simplest way to justify this claim is to assume that each element framework is redundant (one at a time) and then disprove this assump presenting a counterexample. This also proves that the elements o framework are independent (uncorrelated). To do this, we specify a co decision problem (again in the car driving domain since now we are familiar with it) and then change the values of each element of framework, one at a time, and show that this potentially yields a di decision each time. This implies that if that element of the framewo missing, then an ambiguity is allowed in the decision process.

Property: The framework presented in chapter 3 for deciding wheth plan to react to a given contingency in a given situation is of redundancies, i.e. each element of the framework is neces for the final decision.

Justification: we shall state a problem, assume in turn that each elem
it is redundant, and show by counterexample that this is not
Example problem:
Variables:
Situation Space:
Problem: Carry book from home to the office
Plan: Drive car
Context: School time (Weekday, morning, May)
Action: drive straight on Street S, 25 mph
Internal Expectations: reach school
External Expectations: children in sight
$Time_s$: 1 - 3 minutes
Contingency: Ball in front of car
Criticality Space:
$Time_p$ ($Time_{rc}$): very high (very short) (9)
Consequences: small (3)
Side-effects: medium-high (6)
Likelihood: medium (5)
Parameters:
Expert Model:
T_{max} = 9.5; - maximum time pressure allowed to respond
T_{min} = 3; - minimum time pressure required to react
CS_{min} = 4; - maximum difference allowed between side-effects and consequences
L_{min} = 3; - minimum likelihood required to react
MON = 1000; - minimum criticality required to monitor
Agent's Knowledge:
7 contingencies: 4 of higher criticality than this one,
2 of lower criticality than this one
Reactive Planner Model:
decision trees:
f_t: log: 0.2 * $\log_2$(nr_of_conting_with_greater_criticalit
Agent Model:
computational overload - implies computational time delay:
f_{r0}: 1.3 * $time_r$
Behavior Model: "normal"
Parameters for the criticality function f_c: $p_1 > p_5 > p_6 > p_2$
$p_1 = 5$; $p_2 = 1$; $p_3 = 0$; $p_4 = 0$; $p_5 = 3$; $p_6 = 2$

Changing one element of the framework at a time produces following changes in the criticality space values and implicitly in reaction value of this contingency (which imply changes in the of including the contingencies in the reactive plan):
Situation Space:
Problem: Carry 3 kg of radioactive material
Changes: increases Side-effects
Plan: Ride a bike
Changes: increases Consequences
decreases Side-effects
Context: Night-time (non-school time)

Changes: decreases Likelihood
Action: drive straight, 40 mph
Changes: increases Consequences
increases Side-effects
increases Time pressure
Internal Expectations: reach railway crossing
Changes: decreases Likelihood
External Expectations: train in sight
Changes: decreases Likelihood
$Time_s$: ≤ 0.5 seconds
Changes: decreases Likelihood

Note: Any of the changes in the situation space dimensic mentioned also changes the set of possible contingencies which in the one under consideration. Some of the changes add continge with high criticality, so this contingency will get a smaller priorit being considered for reactive response, others have the opposite e

Contingency Child in front of car
Changes: increases Consequences
Expert Model
T_{max}: lower (8.5)
Changes: decreases Criticality (as a whole) since $time_p$ (9) becomes greater than T_{max} (8.5)
T_{min}: higher (9.1)
Changes: decreases Criticality (as a whole) since $time_p$ (9) becomes smaller than T_{min} (9.1)
CS_{min}: lower (2.5)
Changes: decreases Criticality (as a whole) since the difference side_effects - consequences (3) becomes greater than CS_{min} (2.5)
L_{min}: higher (6)
Changes: decreases Criticality (as a whole) since likelihood (5) becomes smaller than L_{min} (6)
MON: higher (i.e. higher than the criticality of this contingency)
Changes: do not even monitor (or prepare to react to) this contingency
Agent's Knowledge:
larger: 24 critical contingencies (more critical than this one)
Changes: the chances to prepare reaction to this contingency decrease because it has a low reaction value compared to the other contingencies known for the same situation
Reactive Planner Model:
decision lists:
f_t = linear: 0.2 * nr_contingencies_with_greater_criticality
Changes: increases real response time
Agent Model:
f_{r0}: 1.8 * $time_r$

5

Changes: increases real response time which may determine it to exceed $time_{rc}$ and thus to be excluded from the reactive plan

Behavior Model: - changes in the criticality function's (f_c) parameters:

p_1: lower (1)

Changes: decreases criticality - disregards allowed response time

p_2: higher (3)

Changes: increases criticality - stresses consequences

p_3: higher (2)

Changes: increases criticality - stresses side-effects

p_4: higher (2)

Changes: increases criticality - stresses anything that can go wrong (both consequences and side-effects)

p_5: lower (1)

Changes: decreases criticality - disregards consequences

p_6: higher (5)

Changes: increases criticality - stresses likelihood (prepares first for the most frequent contingencies)

All these changes in the parameters values of the critical function denote a change in the behavior model implemented b framework, and have as effect a change in the ordering contingencies by reaction value, which may yield a different se contingencies to be selected for reactive response.

o

This concludes our justification that each element of our framewo necessary for the final decision, or alternatively, that the framework is redundancies. We have shown that for any such element, there may variation in its value which may determine a different outcome of the decision, and also that such a variation in this value is possible (and plausible) in the domains under consideration.

5.3. Consistency

I would have liked to be able to say that I proved that the frame introduced in chapter 3 is always able to decide, for any given situation of a set of contingencies possible in that situation should be select planning time to prepare reactive responses for. This would obviously this problem forever, and we could all do something else. But since objective is to design a framework that is applicable in the most dem real-life domains, theoretically proving this property is beyond our However, we are able to theoretically justify a few weaker properties would still ensure the usefulness of the framework. On an encouraging the previous claim actually held in the domains in which experiment conducted. And since these domains are significantly varied in nature, still conclude that it will be true for a large set of real-world domains.

We present here the theoretical justification for our claim that framework for deciding whether to plan to react defined before consistently implement behavior models. This actually means that the o which the contingencies associated with a certain situation are clas according to their reaction value by our framework is the same order a: by the behavior pattern under consideration.

In order to construct our justification, we start with a few prepa definitions and we will prove a few other properties along the way too.

Definition: An Evaluation function (f_e) is a function which, given a set of conditions (pairs contingency-reaction) and a situation in w they apply, computes a score, with the property: the higher score, the better (more appropriate) that set of contingencie according to a particular reaction philosophy (behavior model).

Definition: A Behavior model is an order relationship on the set of contingencies associated with a situation.

The behavior model represents the type of reactive behavior exh by the agent, that is, given any pair of contingencies and their reactior situation, which contingency is to be preferred by the agent for reactic has priority in reacting to, and hence in preparing a reaction for).

Obs.: there is a functional relationship between evaluation function behavior models, i.e. every evaluation function characterize: behavior model, but a behavior model may be characterized by of evaluation functions.

Definition: A Rational behavior is a subset of conditions (pairs contingency-reaction) such that, given an evaluation function an agent with limited resources, there is no other subset conditions that gives a better score for this function while satis the resource limitations.

The notion of rational behavior has been defined independently c situation characteristics, because all the contingencies that belong to th subset must first of all apply to the same situation. The only contributior situation space to the framework is to uniquely define each situation, aı unambiguously identify the contingencies and the reactions associated v

The criticality function f_c (section 3.3.2) defines an order rela called "more important " on a set of conditions matching a given situatic

Definition: Condition a is more important than b in situation S ($a >_S b$) if and only if:

(i) both conditions a and b match situation S

(ii) in situation S: $f_c(a) > f_c(b)$, i.e. the criticality value of a is hi than that of b.

Obs: "more important" is not a partial order relation on the entire conditions in the agent's knowledge base, because there may be situations (S and T) in which both contingencies a and b may appear and such that $a >_S b$, and $b >_T a$. Therefore, the relation "more important" is only defined in a given situation.

Property: The sum of the criticality values (reaction values) for a s conditions is an evaluation function.

Justification: Let $f_c(c)$ be the reaction value of condition (p contingency-reaction) c in situation S, and let C be a set of conditions associated with situation S. Then:

$f_e(C) =)$

is an evaluation function. Indeed, according to the previ definition of an evaluation function, f_c computes a score for a set of conditions in a situation, and since $f_c(c) \geq 0$ for any $c \in C$ (according to its definition in section 3.3.2), f_e can characterize a behavior model. This is true because for any condition a and any set of different conditions C, $C \cup \{a\}$ must be preferred to C by the behavior model (it is never worse to be prepared to react to m contingencies, when agent resource limitations are not taken account, and here the behavior model has been defir independently of the agent's resource limitations).

o

Property: For any two conditions a and b associated with a same situation S, we have: $a >_S b$ if and only if, in situation S, the behavior model prefers condition a to condition b, i.e. it requires the agent to a to include the reactive response for condition a before attempting to include a reactive response for b (i.e., according to the behavior model, given a choice, it is more important that the agent is pre to react to contingency a than to contingency b).

Justification: If $a >_S b$ then $f_c(a) > f_c(b)$ in situation S, so for any s conditions C not including a and b:

$$f_e(C \cup \{a\}) =) + f_c(a) >) + f_c(b) = f_e(C \cup \{b\}),$$

so the evaluation function gives a higher value for $C \cup \{a\}$ than $C \cup \{b\}$, and thus the behavior model requires the agent to attem include a before b in the reactive plan associated with situation S

Conversely, if the behavior model requires the agent to attemp include a before b in the reactive plan associated with situation S, then the evaluation function for this behavior model gives a h value for $C \cup \{a\}$ than for $C \cup \{b\}$, for any set of conditions C applicable in situation S and which do not include a and b, i.e.: $f_e(C \cup \{a\}) > f_e(C \cup \{b\})$, i.e.:

$$) >), \quad \text{that is:}$$

$$) + f_c(a) >) + f_c(b),$$

and so $f_c(a) > f_c(b)$ in situation S, i.e. $a >_S b$.

o

Property: The framework presented in chapter 3 for deciding wheth plan to react to a given contingency in a given situati consistently implements behavior models.

Justification: The notion of behavior model is only about the preferen reacting to contingencies, and thus it is only connected to the r of reaction value, implemented in the framework by the critic function. The previous property shows that the "more impor relation introduced by the criticality function orders contingencies applicable in a situation in the same way as preferences of the behavior model described by this critica function. Therefore, the criticality function represents appropriate way to describe a reactive behavior model in framework.

o

Moreover, because of the optimality property proved in the section, the framework, using the criticality function ordering contingencies, can always optimally implement the behavior mode restricted by the agent's resource limitations, i.e. the rational behavior.

This concludes our justification for the consistency property of framework. This last property has stopped short of claiming that framework is sufficient to simulate any behavior pattern desired, theoretically there are an uncountable number of behavior models and countable number of implementable criticality functions (as a subset of of all programs written in a given programming language), so this woul been impossible to prove (actually we just explained it to be false). Ho we shall state a much more practical conjecture here.

Conjecture: for each known (cited in the literature) type of behavior, exists a combination of parameters in our criticality function w implement it.

This conjecture cannot be actually proved, but can be experime supported, as discussed in section 6.3. Coupled with the previous prop implies that the framework yields the rational behavior for the agent an evaluation function (a behavior model), for any distributions of the characteristics for the conditions (including any distribution of deadli the reactions to contingencies) and for any distributions of the a resources.

If we are unable to come up with a suitable criticality function f desired behavior model, then any of a significant number of automa interactive learning methods may be employed to learn this functic suggested in chapter 7.

5.4. Optimality

We also claim about our framework that it makes the best use of execution time resources of the agent. This means that, given a se contingencies for a situation, the framework will choose not only contingencies that are most important to be treated reactively (accordin reactive model), but will also select as many as it can so that the reacti built for these contingencies maximizes the use of the agent's ru resources.

We first restate here the definition for a rational behavior intro in the previous section:

Definition: A Rational behavior is a subset of conditions (pairs contingency-reaction) such that, given an evaluation function an agent with limited resources, there is no other subset conditions that gives a better score for this function while satis the resource limitations.

According to this definition, a rational behavior maximizes evaluation function for a given situation and agent model, while in the time producing a behavior pattern consistent with the agent's behavior

Property: In the assumptions of section 5.1, an agent, enhanced wit framework presented in chapter 3, exhibits the rational behavi maximizes the use of its resources, while simulating the desi reactive behavior pattern.

Justification: The criticality value establishes an order on the se contingencies associated with the situation, according to the de evaluation function (conf. section 5.3). The decision proc (function f_r in section 3.4.3) is then applied to each of thes contingencies, in the order established. There are two poss outcomes of this process for a contingency which was alre considered worth monitoring: if there are enough resources estimated by the agent model) then the contingency will be inc in the reactive plan; otherwise, this contingency will not included for reactive response. However, this does not mean tha agent's resources were exhausted by the set of contingencies al considered. It only means that the resources left available are sufficient to respond to this contingency (while still respondin useful time to the ones with higher criticality, already include the reactive plan). Therefore, our framework continues evaluation of the remaining contingencies (also in the order their criticality values) since a less critical contingency m require less resources and therefore can also be included reactive response. No contingency can be added to this set w each of the remaining contingencies requires more resources left available by the ones already in the set. Therefore, we con that our framework makes the best use of the agent's resource estimated by the agent model) given a certain evaluation fun (which expresses a specific desired reactive behavior of the age

o

We have thus theoretically justified our claim that the framewor have introduced in chapter 3 for deciding whether to plan to react to contingency in a certain situation yields the rational behavior for the given an evaluation function (a behavior model), for any distributions set of characteristics for the conditions (including any distributio deadlines for the reactions to contingencies) and for any distributions agent's resources. This fact takes off some burden of our experiments, si will only have to conduct experiments for the claims which have theoretical justification. However, we also present, in chapter 6, the res an experimental demonstration of the rational behavior claim as well claims justified in the previous section.

Chapter 6
Experiments

In order to demonstrate the applicability and scalability of the re decision framework presented in chapter 3, we have run a numbe experiments. We describe here these experiments and the main concl that can be drawn from them. In order to demonstrate the generality framework, we have conducted the experiments in three different do the driving domain from which we took most of the examples used duri previous presentations, and two medical domains of expertise: anesthe and intensive care patient monitoring. It is well known that different in a domain may have different opinions on specific subjects from the In order to obtain a consensus of these opinions in the driving doma have polled 8 experts, and we have combined their opinions in differen It was interesting to find out that the results of these combinations had degree of similarity among them, and were well in line with the indi opinions of the experts (although among them opinions may have significantly). For the medical domains we have only used the advice single expert in the field. In the following section we describe the know acquisition process which we have conducted in the driving domain, a results. Then we describe a set of experiments in this domain, that supp claim of optimality for our framework which has been theoretically ju in the previous chapter. In the third section we present a set of expe which were aimed to demonstrate how different behavior models ca described in our framework and how they affect the reactive behavior agent using them. We conclude this chapter with a description of ho reaction decision framework proposed here can be included in a co agent which runs in a real, complex world.

6.1. The Driving Domain

In this domain we were able to collect knowledge from 8 experts, most people can be considered experts in this domain, and seven c colleagues (David Ash, Alex Brousilovski, Lee Brownston, Janet Murd Serdar Uckun, Rich Washington and Michael Wolverton) were kind enou volunteer their valuable time and experience for this part of the p Beside providing the raw knowledge, they have also made signifi comments which have helped me clarify the knowledge acquisition pr involved. I am indebted to all of them (the eighth person in the expe was myself).

Contingency
Reaction

1	Child runs from right, 20 m in front of car	Brake hard and steer right
2	Car crosses w/o priority 20 m in front, from right to left	Brake and gently steer right
3	Car in front stops suddenly	Brake hard
4	Cat runs across street, 20 m in front	Brake hard and steer right gently
5	Traffic light changes red 40 m in front	Brake hard
6	Tire explosion	Brake gently and do not steer
7	A deep and medium width hole detected 30 m in front	Brake hard and steer right gently
8	Airplane lands in front of car	Brake moderately hard
9	Brake malfunction light turns on	Brake gently
10	Engine overheat light turns on	Brake gently to stop the car
11	Loud radio turns on suddenly	Adjust radio volume
12	Meteor falls on the trunk of the car	Accelerate hard
13	A ball pops in the street, from the right, at 20 m in front[5]	Brake hard and steer right

Table 6.1. Contingencies for the car driving domain experiments

Table 6.1 lists the 13 contingencies (also listed in table 3.1 and use illustration purposes throughout the thesis) proposed, together wit reactions for each of them.

The knowledge acquisition problem was to specify a value betwee and 10 for three of the criticality space dimensions (consequences, side and likelihood), and a real time value for the time to respond tc contingency, for each of these 13 contingencies, when considered possi appear in the following situation:

Problem: Deliver package to work
Plan: Drive car
Context: May, midweek, morning (school time), pass in front of a sch
Ext. Expect.: Children in sight
Int. Expect.: Reaching school zone
Action: Drive straight, 25mph
$Time_S$: max. 3 minutes

The experts were instructed to translate their qualitative feelings quantitative values, and to concentrate more on relative values than absolute values they were giving. As some of them have commented, the used was sometimes closer to logarithmic and sometimes closer to expor but very seldom (if ever) was it approximately linear.

Each expert was also independently asked to order the set contingencies by reaction value, that is, to specify the order in which she believes the agent should consider these contingencies for reactic well as where a threshold on monitoring for them should be placed. information was then used to evaluate the results of applying our frar to the data supplied by the experts. The experts were asked to provic contingency dependent knowledge independently of the final orderi any case, we believe there is little danger of any conscious correl between the data supplied by an expert for each contingency and the preference specified by the same expert, because of the amount of info they had to supply - over 50 values each only for contingency data.

I will omit here the individual values supplied by each expert for contingency precisely because of the considerable amount of nu involved. I would prefer to comment instead a little on the distribution c values, although a meaningful statistical analysis would not be fully re because the still small number of experts involved. The absolute v specified varied quite a lot. For example, the consequences of not react the engine overheating was rated between 4 and 10 (on the scale of 0 where 0 meant no consequences at all), while the likelihood of a running in front of the car was rated between 4 and 9. Although the c of the contingencies differed too (traffic light was placed between firs seventh while airplane, radio and meteor all varied between the ninth a thirteenth places), the experts opinions agreed much more on the s contingencies to be actually taken into account (i.e. the monitc threshold): all of them indicated the first four contingencies as orde table 6.1, all but one indicated the hole, and all but two indicated the the tire contingencies.

This was the first indication that, although individual pairs of e may disagree, each of the experts tends to agree with the opinion of the This conjecture was then supported by a deeper analysis of the rest of t supplied by the experts. We have further analyzed the order specified experts on the set of contingencies, by assigning an order number tc contingency according to each expert's specification, and then taking median value, average value, and average of the set of 6 numbers obtai eliminating the highest and lowest expert specified value for (contingency independently. In all three cases, the result of orderin contingencies by the values obtained this way were identical, and differences with each expert were much smaller than differences be individual experts. This again supports the previous conjecture. It wa: interesting to see that not even one expert had specified the same orde inferred by all the three statistical methods. A further confirmation conjecture came from the fact that, for each characteristic of contingencies, the three statistical measures have produced very s results. Moreover, after eliminating the two extreme values in each cas remaining values were much closer, which shows that the experts te: agree with each other most of the time. Also, since different experts different scales to measure the same qualitative phenomenon, the qua aspects of their input (orderings) tend to agree more than the quan formulations (the experiments described further in this chapter will that our framework is robust to quantitative variations in the knov specification, and is well suited to extract the qualitative aspects of it, are the ones which eventually interest us).

The analysis of this data also suggested that different experts tak same (or consistent) decisions, but apparently for different reasons, t they have different heuristic "formulae" or rules to combine their eva of the characteristics of events in their domain of expertise. All t observations support our explicit inclusion in the framework of an model, which has the role of calibrating the entire reaction dec framework according to the set of qualitative_to_quantitative transfoi functions used by the expert providing the domain knowledge.

Contingency

Time$_{rc}$
Time$_{p}$
Consequences
Side-effects
Likelihood

1
Child
1.0
10.0
10.0
6.5
4.5

2
Car-X

5

1.0
10.0
8.8
5.7
4.0

3
Car stop
2.0
5.0
7.0
3.2
6.2

4
Cat
1.0
10.0
5.2
5.9
6.8

5
T.light
4.0
2.5
6.2
0.7
8.8

6
Tire
3.0
3.3
6.0
2.8
2.3

7
Hole
2.0
5.0
4.5
4.5
2.8

6

8
Plane
2.5
4.0
9.5
4.5
0.3

9
Brake
30.0
0.3
6.2
1.0
2.0

10
Heat
50.0
0.2
5.5
0.3
2.2

11
Radio
100.0
0.1
1.8
0.7
2.0

12
Meteor
0.1
100.0
9.5
3.2
0.1

13
Ball
1.0
10.0
0.7

7

5.7
5.0

Table 6.2. Data values for the car driving domain experiments

In our experiments conducted with data from the driving domair have used the average_after_extremes_elimination values, obtained frc raw data provided by the experts as described above. These values presented in table 6.2. The order in which the contingencies are preser both tables 6.1 and 6.2 is the average_after_extremes_elimination (wh mentioned above, is the same as the average and the median) order o from the pool of experts. The experiments with this set of data are l presented in the next two sections.

6.2. Optimality

We present here the results of the experiments we have conducte support the theoretical claims made in chapter 5. Since most of these were justified theoretically, these experiments are merely demonstrati applying the framework. We have used four different reactive planner and five agent models to show how the recommendations of the fram vary and how it continues to ensure the optimal use of the agent's re for the given agent models.

We have also used a "normal" behavior model, that is we expect agent to behave the same way as the experts recommend. In calculatir reaction value of a contingency, this model assigns more weight to the pressure dimension, followed by the difference between consequence side-effects, and then likelihood. Consequences are taken into account b themselves, but also (and mostly) in combination with the side-effects. the criticality function parameters given by the behavior model are:

$$p_1 = 5,\ p_2 = 1,\ p_3 = 0,\ p_4 = 0,\ p_5 = 3,\ p_6 = 2,$$

where the parameters specified by the expert model (an abst "average_after_extremes_elimination" expert) are:

$$T_{max} = 20.0;\quad T_{min} = 1.0;\quad CS_{min} = 2.3;\quad L_{min} = 1.3;\quad MON = 10000,$$

and the function computing the time pressure is:

$$f_{tc} = 10 \,/\, time_{rc}\ .$$

In this particular case, the criticality function (described in se 3.3.2) becomes:

Criticality = f_c (t, c, s, l) =

if	(t > 20)	then $f_c = 0$
elseif	(c + 2.3 - s < 0)	then $f_c = 0$
elseif	(t < 1)	then $f_c = {}^3 * l^2)$
elseif	(l < 1.3)	then $f_c = {}^3 * l^2)$
else		$f_c = t^5 * c * (c+CS_{min}-s)^3 * l^2$

Table 6.3 presents the values returned by this function, and the ou of the monitoring decision of the framework. We provide them only to the reader to 'feel' the results of the framework. The monitoring thresh set by the expert model in a region of the contingency space where the substantial gap among the reaction values of the contingencies order criticality. Since the expert and behavior models do not change durir

experiments described in this section, these values will not change e
They will however change anytime at least one of the expert or beh
models change. In the experiments described in the following sections w
not include this criticality value anymore. It can however be ea
recomputed from the behavior models, which will always be specified.

	Contingency	Criticality	Monitor
1	Child	3.95E9	yes
2	Car-X	2.21E9	yes
3	Car stop	1.90E8	yes
4	Cat	9.84E7	yes
5	Traffic light	2.22E7	yes
6	Tire	2.17E6	yes
7	Hole		

9

1.34E6
yes

8	Plane	5.83E2
9	Brake	6.56
10	Heat	1.89
11	Radio	5.3E-2
12	Meteor	0.00
13	Ball	0.00

Table 6.3. Criticality values for the "normal" behavior model, for the car driving domain experiments

The first and most important observation of the experiment is tha framework orders the contingencies by criticality value (based on the from the "average" expert) identically to the order indicated by the "average" expert. When presented with this ordering, all the human involved have agreed to its rationality.

We must also point out here that the framework proved very robu that considerable variations in the values of the behavior and expert parameters as well as in the absolute values for the dimensions contingencies have yielded the same order induced by the criticality fu What really matters is the relative relationship among pairs of elements framework. For example, in the normal behavior model, time pressur greatest weight. We have experimented with variations of up to 25% i absolute value of its weight (p_1) and have still obtained the same ord have repeated the experiment for other behavior models were time pres also considered most important, as well as by varying other paramet slightly varying individual values of the characteristics of contingencie in each case we have obtained robust behaviors of the framework. suggests that small variations in the values provided by experts shou negatively influence the behavior of an agent using this framework.

In the experiments described in this section, we have used the foll four reactive planner models:

RP1: constructs balanced binary decision trees; the function estimatir global reacting response time:
$ft = k_p * \log_2$ (number_of_contingencies_with_≥_criticality),
where the average test time is: $k_p = 0.2$ seconds.

RP2: same as RP1, but the average test time is: $k_p = 0.3$ seconds.

RP3: constructs decision lists; the function estimating the global rea response time is linear:
$ft = k_p *$ number_of_contingencies_with_≥_criticality,
where the average test time is: $k_p = 0.2$ seconds, and the decisior are built such that the pre-conditions discriminating t contingencies with the highest time pressure are tested first.

RP4: same as RP3, but the average test time is: $k_p = 0.3$ seconds.

We have also used five agent models. The only difference among the the computational load estimated to be imposed on the agent at executi (for this situation), which has the effect of slowing the agent, that increases the response time of the agent to a contingency by a factor K_t:

$$f_{r0}(time_r) = time_r * K_t ;$$

The five agent models used are:

AM1: $K_t = 1$, that is, there is no computational overhead estimated;

AM2: $K_t = 1.3$, that is, there is a 30% computational overhead estimated;

AM3: $K_t = 1.8$, that is, there is an 80% computational overhead estimate

AM4: $K_t = 2.5$;

AM5: $K_t = 4.0$.

Contingency

Monitor

React (RPModel = decision trees - $k_p = 0.2$)

$K_t = 1.0$
$K_t = 1.3$
$K_t = 1.8$
$K_t = 2.5$
$K_t = 4.0$

1
Child
yes
yes
yes
yes
yes
yes

2
Car-X
yes
yes
yes
yes
yes
yes

3
Car stop
yes
yes
yes
yes
yes

4
Cat
yes
yes
yes
yes
yes

5
Traffic light
yes

12

yes
yes
yes

6
Tire
yes
yes
yes
yes

7
Hole
yes
yes
yes

8
Plane

9
Brake

10
Heat

13

11
Radio

12
Meteor

13
Ball

Table 6.4. Optimality demonstrations results for reactive planner model

Contingency

Monitor
React (RPModel = decision trees - k_p = 0.3)

K_t = 1.0

K_t = 1.3
K_t = 1.8
K_t = 2.5
K_t = 4.0

1
Child
yes
yes
yes
yes
yes
yes

2
Car-X
yes
yes
yes
yes
yes

3
Car stop
yes
yes
yes
yes

4
Cat
yes
yes
yes

5
Traffic light
yes
yes
yes

6
Tire
yes
yes

7
Hole
yes
yes

8
Plane

9
Brake

10
Heat

16

11
Radio

12
Meteor

13
Ball

Table 6.5. Optimality demonstrations results for reactive planner model

Contingency

Monitor
React (RPModel = decision lists - k_p = 0.2)

K_t = 1.0
K_t = 1.3
K_t = 1.8
K_t = 2.5
K_t = 4.0

1
Child
yes
yes
yes
yes
yes
yes

2
Car-X
yes
yes
yes
yes
yes

3
Car stop
yes
yes
yes
yes
yes
yes

4
Cat
yes
yes
yes

5
Traffic light
yes
yes
yes
yes
yes
yes

6
Tire
yes
yes
yes
yes
yes
yes

7
Hole
yes
yes
yes
yes
yes

8
Plane

19

9	Brake
10	Heat
11	Radio
12	Meteor
13	Ball

Table 6.6. Optimality demonstrations results for reactive planner model

Contingency

Monitor
React (RPModel = decision lists - k_p = 0.3)

K_t = 1.0
K_t = 1.3
K_t = 1.8
K_t = 2.5
K_t = 4.0

1
Child
yes
yes
yes
yes
yes

2
Car-X
yes
yes
yes

3
Car stop
yes
yes
yes
yes
yes

21

yes

4
Cat
yes
yes

5
Traffic light
yes
yes
yes
yes
yes
yes

6
Tire
yes
yes
yes
yes
yes
yes

7
Hole
yes
yes
yes
yes

8
Plane

9
Brake

10
Heat

11
Radio

12
Meteor

13
Ball

Table 6.7. Optimality demonstrations results for reactive planner model

Tables 6.4 to 6.7 summarize the results of our demonstrations. The
the set of contingencies recommended by our framework for rea
response preparation, in each case. As expected, this set decreases wi
increase in the agent computational load, all other things being
(different columns in the same table). It also decreases with an increase
cost (here time) of the average tests to be performed (as can be se
comparing the corresponding columns in tables 6.4 and 6.5, as well as th
columns in tables 6.6 and 6.7. In each case, the agent tries to optimize
of the agent resources (i.e. to include as many contingencies as pos
while maximizing the evaluation function on the subset of sele
contingencies, by essentially including the highest criticality conting
possible. Obviously, the more accurate the agent and planner models us
better the selected contingencies will actually optimize the use of ru
resources (the models used here are quite rough - assume all tests ta
same time and that the simple logarithmic and linear formulae stated
correctly approximate the agent).

In this example the decision trees model always selects t
contingencies in the strict order of criticality (which need not be the
general), while the decision lists model allows for gaps in the strict ord
that it can accommodate a larger number of contingencies. This is one
proof that the algorithm proposed in chapter 3 optimizes the use of the
resources. For example, in table 6.7, for $\tau = 2.5$K, the agent can respond to only
one contingency with a response time of maximum 1 second, so it choos
one with largest criticality (the child contingency); it can respond to
contingencies with maximum response time of 2 seconds (the car_stop
hole contingencies), and so on, but cannot respond to the other contin
with short (1 second) response time, so it will omit them from the fin
Also note that the decision lists based planner model assumes tha
contingencies are ordered by the response time allowed (in the final r
plan), and also that the test times for each contingency are constant.
first of these assumptions would have not been included in the re
planner model, then the default assumption is that the contingenci
ordered by criticality, and then the reactive plan for this case could n
included the hole contingency since it would have been last in the d
list, and its response time would have exceeded its allowed response time.

One last observation from these experiments is that, for this part
set of data, it confirms our discussion of decision trees versus decisio
from section 3.4.1. We argued there that there are frequent cases in wh
set of contingencies recommended by the framework is larger when us
decision lists planner than a decision trees planner, all other things
equal (which may seem somewhat counterintuitive at the first glance).
in this demonstration, the decision lists based agent includes
contingencies than its decision trees based counterpart for most of the
covered. In our example, the evaluation function value is usually grea
the decision trees case, because of a subtle violation of the "all other
being equal" assumption: the decision trees based planner model assum
there is no test time needed to reach a response for a single contin
($\log_2 1 = 0$), while the decision lists based planner model assumes that t
needed to reach such a response is still the time needed to perform one
this assumption would have been made in the first case too, then the

lists planner model would have yielded also a higher evaluation functio its corresponding decision trees counterpart, for the set of conting recommended in at least some cases (like RP2 and RP4 (k_p = 0.3) and AM4 4.0)).

6.3. Behavior Models

Though not intended as a simulation of human behavior, our app to solving the reaction planning decision problem has some pote applications in this area too. Specifically, it provides the basis for a p language to discuss the characteristics of different human behavior related to this task. In this section we shall propose a way of represent our framework some such behavior models discussed in the literature, a as the results of a few experiments we have done using this represen Our discussion here is by no means intended to give a complete solution problem of simulating human reactive behaviors, but is only intend suggest a possible such representation, which needs a lot more resea prove its usefulness or to find its best application domain.

In section 5.3, we have justified the property that our reaction de framework consistently implements behavior models. We stated ther conjecture that for most types of reaction-related behaviors cited i literature, there is a corresponding behavior model encoding in framework which implements that type of reaction. Here, we go even a further, by defining a couple more such behavior models and repres them in our framework too. Since we found no way to theoretically prov conjecture, we have conducted a number of experiments designed to sup which we present in this section. They show how our framework determine an agent to exhibit different reactive behaviors for the domain described before, while also helping us to clarify the meaning different thresholds and parameters in our framework.

Besides the so called "recommended" or "normal" behavior, we found six more types of reaction-related behaviors - sometimes hazardous attitudes [Woods & al., 1987; FAA, 1991]. The last two behavior proposed by David Gaba (personal communication, 1993). Here is a description of each of these behaviors:

- Recommended Behavior - is the normal behavior expected by the exper and from an expert in the domain.
- Antiauthority Behavior - is the "don't tell me!" type of behavio which the agent regards rules, regulations and procedures unnecessary, and thus tends to disobey them.
- Impulsivity Behavior - is the "do anything quickly!" type of behavior, i which the agent attempts to always do the first thing that come mind, without stopping to think and select the best alternative.
- Invulnerability Behavior - is the "it won't happen to me" type behavior, in which the agent is always inclined to take risks sinc believes that the current situation is never one of those (less likel still possible) situations when something wrong might just happen
- Macho Behavior - is the "I can do it!" type of behavior, in which agent wants to impress others, and is ready to take significant risks it. It is inclined to react even when not really necessary or when it be more dangerous than not to react. Such agents either forget

the possible side-effects of their actions, or at least discount de
these side-effects.

m Resignation Behavior - is the "what's the use?" type of behavior, in
which the agent faced with a critical situation usually chooses to
nothing, since it underestimates its capacity to respond to the even
the effectiveness of such a response, in the given time frame. It h
tendency to leave such actions to others, for better or for worse.

m Risk-averse Behavior - the agent tries to avoid risk by all means
(considering both the consequences of not being prepared to rea
time, and the possible side-effects of reactions), but may therefore
sometimes less importance to the time pressure.

m Liability Conscious Behavior - the agent is particularly interested in
avoiding any legal liabilities that may arise from its actions. There
it tends to prepare to always do something, preferably what is le
bounding, even if that something may be believed not to succeed in
particular situation. This may prevent the agent to prepare for
other contingencies which are less liability creating, but which c
have been treated if there were enough resources available.

m Social Responsibility Behavior - the "socially conscious" agent tend
put the interests of the society before those of the individual, inc
itself.

Each of these behaviors can be simulated in our framework by adj
the parameters of the corresponding behavior model. While the a
parameter values are less important, their relative values define the d
behavior models.

	Behavior Model						Expert Model			
Behavior	p_1	p_2	p_3	p_4	p_5	p_6	T_{max}	T_{min}	CS_{max}	L_{min}
Recommended	5	1	0	0	3	2	20.0	1.0	2.3	1.3
Antiauthority	5	1	0	0	3	0				
Impulsivity	0	0	0	0	0	3		10.0		5.0

Invulnerability
5 1 0 0 3 2
4.2

Macho
4 1 0 0 0 3
10.0

Resignation
5 1 0 0 3 2
5.0

Risk-averse
2 2 4 2 1 1

Liability conscious
3 3 1 2 1 2
100.0 0.0 0.0

Social responsibility
4 3 0 0 4 3

Table 6.8 Representing Behavior Models

Table 6.8 summarizes the representation of these behavior models our framework. Recall that a behavior model in our framework is imple by a set of values for the parameters of the criticality function (comput reaction value of the contingencies), and may also be influenced by values of the thresholds given by the expert model (section 3.3). The val the expert model parameters are completely specified in the table only recommended behavior model; for the other models, only values that changed from the initial specification are given. Also remember that pressure is the only parameter which can take values outside the in [0,10], because it is converted from arbitrary real values using the specified conversion function[10] Therefore, the time pressure related parameters are harder to generalize among domains, as will be notic appendix 3, where we present the results of the same experiments run anesthesiology domain, with the same parameter values as here except time pressure dimension. The expert models in table 6.8 were used i demonstrations in the driving domain.

To illustrate the simulation of these behaviors in our framework have run the framework with the behavior models presented in table the 13 contingencies presented previously the driving domain. Tabl summarizes the results of these experiments. We have also shown the r values produced by the criticality function. Their absolute values ha meaning whatsoever; what matters are their relative values (and only the same behavior model), which represent the relative value of react one contingency vs. another in a same situation and under the same b model. For each behavior, the monitoring threshold was set (throug expert model) in a region of the contingency space where there substantial gap among the reaction values of the contingencies order criticality. The threshold is represented by a thicker line separatin contingencies for each behavior into two sets. The numbering contingencies for each behavior model is the same as for the recomm behavior. This was done in order to facilitate comparisons of each be model with the "normal" one.

In chapter 5, we have defined a behavior model to be an o relationship on the set of contingencies associated with a situation. The in the experiments described in this section, we only concentrate o ordering of contingencies by reaction value (and sometimes relative val the criticality function, but never on its absolute values), and ignor issues related to the reactive planner model and the agent model, that ignore the final decision of applying the framework to a set of conting This is consistent with the purpose of our demonstrations here, sinc specific agent (with a given reactive planner and resource limitations exhibit any of the reaction behaviors discussed, depending only on the in which its behavior model recommends the contingencies for consid to be reacted to, and not on the actual components and resources of the

Behavior Model 1

Behavior Model 2

Behavior Model 3

(Recommended)

(Antiauthority)

(Impulsivity)

28

1
Child
3.95E9
1
Child
1.95E8
4
Cat
3.14E2

2
Car-X
2.21E9
2
Car-X
1.38E8
5
T.light
26.10

3
Car stop
1.90E8
3
Car stop
4.96E6
3
Car stop
15.43

4
Cat
9.84E7
4
Cat
2.12E6
1
Child
9.54

6
T.light
2.22E7
6
Tire
4.10E5
2

Car-X
8.00

5
Tire
2.17E6
5
T.light
2.87E5
7
Hole
4.68

7
Hole
1.34E6
7
Hole
1.71E5
6
Tire
3.48

8
Plane
5.83E2
8
Plane
1.94E3
10
Heat
3.26

9
Brake
6.56
9
Brake
3.28
9
Brake
2.82

10
Heat

30

1.89
10
Heat
0.86
11
Radio
2.82

11
Radio
5.3E-2
11
Radio
2.6E-2
8
Plane
0.16

12
Meteor
0.00
12
Meteor
0.00
12
Meteor
0.00

13
Ball
0.00
13
Ball
0.00
13
Ball
0.00

Table 6.9 Reactive Behavior Experiments for the Driving Domain

Behavior Model 4

Behavior Model 5

Behavior Model 6

(Invulnerability)

(Macho)

(Resignation)

1
Child
3.95E9
4
Cat
1.63E7
3
Car stop
1.90E8

3
Car stop
1.90E8
1
Child
9.11E6
5
T.light
2.22E7

4
Cat
9.84E7
2
Car-X
5.63E6
6
Tire
2.17E6

32

5
T.light
2.22E7
3
Car stop
1.04E6
7
Hole
1.34E6

2
Car-X
4.70E4
13
Ball
8.75E5
8
Plane
5.83E2

6
Tire
1.47E3
5
T.light
1.65E5
9
Brake
6.56

7
Hole
1.15E3
7
Hole
6.17E4
10
Heat
1.89

8
Plane
5.83E2
6
Tire
9.01E3
11

Radio
5.3E-2

9
Brake
6.56
8
Plane
8.10
1
Child
0.00

10
Heat
1.89
9
Brake
0.78
12
Meteor
0.00

11
Radio
5.3E-2
10
Heat
0.30
2
Car-X
0.00

12
Meteor
0.00
11
Radio
3.7E-2
4
Cat
0.00

13
Ball

34

0.00
12
Meteor
0.00
13
Ball
0.00

Table 6.9 Reactive Behavior Experiments for the Driving Domain (contir

Behavior Model 7

Behavior Model 8

Behavior Model 9

(Risk-averse)

(Liability conscious)

(Social responsibility)

1
Child
1.2E11
1
Child
2.0E11
1
Child
1.0E12

2
Car-X
3.7E10
2

Car-X
7.0E10
2
Car-X
3.7E11

4
Cat
4.39E9
12
Meteor
3.8E10
3
Car stop
7.0E10

3
Car stop
5.05E8
4
Cat
7.56E9
5
T.light
2.3E10

7
Hole
1.08E8
3
Car stop
3.34E9
4
Cat
2.89E9

6
Tire
2.40E7
5
T.light
7.49E7
6
Tire
2.96E8

5
T.light
1.88E5
7
Hole
7.48E7
7
Hole
3.49E7

8
Plane
1.59E5
6
Tire
5.04E7
8
Plane
4.10E3

9
Brake
57.62
8
Plane
3.17E7
9
Brake
2.72E2

10
Heat
2.33
9
Brake
1.37E4
10
Heat
94.70

11
Radio
0.57
10
Heat
4.87E2

11
Radio
0.78

12
Meteor
0.00
11
Radio
0.34
12
Meteor
0.00

13
Ball
0.00
13
Ball
0.00
13
Ball
0.00

Table 6.9 Reactive Behavior Experiments for the Driving Domain (contir

Here is a brief explanation of the changes required by the paran for each behavior model, with respect to the normal behavior model de in the previous section, as well as the main effects they have on the o of the 13 contingencies we have presented in the previous section, for driving domain:

- Antiauthority Behavior Model - do not take likelihood into account likely events are usually covered by laws, regulations and proced The traffic light contingency goes down in criticality, as the only regulation to be observed as a contingency in our set; the rest re the same.
- Impulsivity Behavior Model - consider a single response, for a contingency with great (but serviceable) time pressure and likelihood, to allow at least for a reasonable response in a signif number of cases; the reactive plan will consist of a single reactio this contingency. Consequences and side-effects are disregarded, time pressure is considered only through raising T_{min} (to 10) so as to include only the high but still acceptable time pressures. Likelihoo the only one still considered in the reaction value formula, and L_{min} is also raised significantly (L_{min} = 5). Therefore, the car contingency becomes the only one selected for reaction preparation.
- Invulnerability Behavior Model - Low and medium likelihood contingencies are considered much less critical ("it won't happe me..."); only high likelihood contingencies are really conside

critical, so the L_{min} threshold is significantly increased (L_{min} = 4.: our tests, the car crossing contingency falls a lot because its likelihood becomes lower than this threshold.

- Macho Behavior Model - Forget about side-effects, and also ta consequences less into account, since the agent mainly tries to im others, by preparing for time pressured, but especially lil contingencies, so that it can react most of the time. The likelil weight is increased, while the CS_{min} threshold is also increased (CS_{min} = 10.0) such that it becomes useless. In our demonstration, ball advances all the way to number 4 because the difference between consequ and side-effects is not considered here, while cat advances to number one since it is more likely than the first three, and its side-effect also disregarded.
- Resignation Behavior Model - Here it is interpreted as underconfidence, that is underestimating its own abilities, since we only talk al reaction preparation at planning time, and not reaction behavi execution time (were it would have been interpreted as 'giving up') agent is willing to prepare to respond only to low time pressured e and therefore the T_{max} threshold is significantly reduced (by 7. T_{max} = 5). Therefore, many contingencies with higher time pres get zero reaction value and fall at the end of the list.
- Risk-averse Behavior Model - taking most precautions to avoid risk, decision process considers mostly the side-effects of the react followed by the consequences of not reacting and the sum consequences and side-effects, and much less time pressure likelihood. The driving domain contingencies become roughly ord by this sum, with a few exceptions: the plane contingency has very likelihood, the brake contingency has very low time pressure, and meteor contingency has a too short response time allowed for a re to be effective.
- Liability Conscious Behavior Model - while the weight of time pres and of the difference between contingencies and side-effects decr the agent assigns more importance to consequences, side-effects their sum. Also, there are no threshold for either time pressur likelihood (T_{min} = L_{min} = 0; T_{max} = 100), since a contingency sh never be discarded only because a reaction to it is believed to be u Therefore, meteor becomes very high priority here, and the agent v prepare roughly in the order of collision with people, moving animals, objects. Ball is still not considered here at all because the side-effects are still much higher (and potentially more liable) tha consequences. Also in this case, more contingencies are considered monitoring than usually.
- Social Responsibility Behavior Model - preparing a population optimal behavior involves considering both consequences alone and difference between consequences and side-effects, as well as likelihood, more than before (with respect to time pressure). It is c to the "normal" behavior described in the previous section, with only difference that traffic light gains priority with respect to since this behavior tends to favor groups of people over single pe and people over animals. Notice here that significant overall cha

in the values of the parameters, but small changes in their rel order, have produced a very similar ordering of the contingenci compared to the recommended behavior.

As can be noticed from the above discussion, the results of t demonstrations require a certain amount of interpretation. This is ne especially since the definitions of these behavior models are generally on execution time types of reactions, while we attempt here to implemen at planning time. However, their interpretation shows that they reasonable and consistent with the generally accepted (execution- definition of each behavior model, and that there is a plausible explana the results that maps them into the corresponding (conceptual) beh These demonstrations show that our framework may at least provi reasonable basis for representing and exchanging information and about reaction-related behavior models, and thus for interpreting studying different behaviors. For example, given a specific behavior (or the set of contingencies), we can automatically discover the parameters behavior model which emulates it, and then we can characterize this b and maybe attempt to correct it.

The specific values of the different parameters of the behavior m used may vary in certain limits, producing essentially the same results fact contributes to the robustness of our framework, and simplifie knowledge acquisition process by easing the burden of specifying ac values for the criticality space dimensions by the expert. More importa the relative values the expert supplies, but this is generally easier to a Also the expert model may influence some of the behavior models, s expert should probably be informed in advance about the desired be model. However, our experiments were conducted without informing expert on the type of behavior model desired, and as can be seen fro discussion here (and also according to our experts), the results a agreement with the definition of each behavior model.

We have also run the same demonstrations on a set of contingenci a situation in the anesthesiology domain. Again the results satisfied the interpretation of the different behavior models. A brief description o experiment and a short interpretation of the results for each behavior are presented in appendix 3.

In the next section we present a final experiment, aimed demonstrating that the framework defined in this thesis can scale up integrated in complex autonomous agents, designed to work in real, co domains, and that by doing this, we improve the agent's global real performance (by making it more responsive to those events that considered more critical in the domain). This way we not only impro quantitative performance of the agent, but more importantly, the quality performance. The experiment presented in the next section was also ai supporting evidence that the knowledge required to apply our fran exists in real domains, that it can be reasonably quantified by experts domain, and that it can be acquired from these experts and produce re results.

6.4. Complex Real World Domain

We present here one more experiment we have conducted with framework, in a real life medical domain: patient monitoring in an in care unit (ICU). This time, our framework was integrated in a com real-time agent architecture capable of planning, reaction, and dy replanning: the Guardian system [Hayes-Roth & al., 1992, Hayes-Roth, Our framework has the role of filtering the information which flows fro planner to the reactive planner, according to the architectural outlined in appendix 1.

The two domain experts who have generously advised us (David and Serdar Uckun) have identified 68 contingencies for a set of situ corresponding to a general intensive care monitoring case (figure 6.1) have also specified heuristic values for the four characteristics for ea these contingencies. For an easier understanding of the presentation, w present part of these experiments and most of the data concerned, in a 4, and shall discuss here only the main results.

Problem:	Intensive care monitoring
Plan:	normal postoperative procedure
Context:	after coronary artery bypass grafting (CABG) procedure 50 years old patient, no other history known
Action:	ventilate patient / weaning / extubate patient
Internal Expect:	
External Expect:	
Time:	0-8 hours / 9-18 hours / 18-48 hours

Figure 6.1. Situations for the ICU domain

Table A4.1 lists the entire set of contingencies and the characte values for them, in the order specified by the experts (grouped by cat of complications that may develop).

The first part of this demonstration consisted in running the crit function part of the framework on this data set, for the recomm behavior model (section 6.3), for several expert models. We have exemplified for a large real-life case, the influence of varying diff expert model parameters, over the ordering of contingencies by criti Appendix 4 presents a partial set of results from this demonstration A4.2 to A4.5).

The most important conclusion to be drawn from this demonstrati that the recommendations of our framework are reasonable from the point of view. Our experts have agreed, in each case (i.e. for each expert used) with the ordering of the contingencies proposed by our system, them reasonable and finding reasonable interpretations for them. Sinc is no other (objective) way to evaluate the framework's recommendatio may conclude that the framework and the "normal" behavior model w defined are a reasonable solution to our original problem.

#	Contingency (Response would be the typical response for this event)	Resp. time	Conse- quences	Side- eff.	Likeli-

41

hood
Criti-
cality

34
et-tube-disconnection
2
10
2
4
4.2E12

18
ventricular-tachycardia
1
9
7
2
2.2E12

13
ventricular-fibrillation
1
10
8
1
6.1E11

35
kinked-et-tube
5
8
2
4
1.8E10

20
hypoxia
5
8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

42

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

1
myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism

43

10
8.5
7.5
3
2.13E7

6
hypovolemia
20
7
3
7
2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax
10
8
7
3
2.01E7

40
acute-hemolytic-transfusion-react
10
8.5
5
1
1.28E7

26
hemothorax
10
7
7
4
1.05E7

9
right-heart-failure
10
8
7

44

2
8.94E6

11
postop-hypertension
20
6.5
5
4
1.38E6

Table 6.10. Selected Contingencies for k_p = 0.5 (30 seconds) for ExplorerII (k_t = 1.166)

The second part of the demonstration considers the behavior of framework in the context of the Guardian system. The blackboard- [Hayes-Roth, 1985] Guardian agent has a reactive planner (ReAct) action-based hierarchies [Ash, 1994].

#	Contingency (Response would be the typical response for this event)	Resp. time	Conse- quences	Side- eff.	Likeli- hood	Criti- cality
34	et-tube-disconnection	2	10	2	4	4.2E12
18	ventricular-tachycardia	1	9	7	2	2.2E12
13	ventricular-fibrillation	1	10			

45

8
1
6.1E11

35
kinked-et-tube

5
8
2
4
1.8E10

20
hypoxia

5
8
6
4
2.53E9

7
myocardial-ischemia

5
8
6
3
1.42E9

15
sinus-bradycardia

5
7
5
3
1.24E9

14
ventricular-ectopy

5
7
7
6
7.62E8

5
cardiac-tamponade

5
8.5
7.5
3
6.84E8

46

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

1
myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism
10
8.5
7.5
3
2.13E7

6
hypovolemia
20
7
3
7
2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax

47

10
8
7
3
2.01E7

40
acute-hemolytic-transfusion-react
10
8.5
5
1
1.28E7

26
hemothorax
10
7
7
4
1.05E7

9
right-heart-failure
10
8
7
2
8.94E6

11
postop-hypertension
20
6.5
5
4
1.38E6

4
increased-afterload
20
6.5
5
4
1.38E6

36
right-mainstem-intubation
20
6.5
3

48

2
1.23E6

16
atrial-fibrillation
20
7
6
4
9.78E5

41
febrile-nonhemolytic-transfus-react
20
6.5
4
2
6.98E5

67
low-k
30
7.5
5
5
6.63E5

42
mechanical-bleeding
20
7.5
7.5
4
3.54E5

66
dilutional-low-na
30
7
2
2
3.48E5

64
low-na
30
7
2
2
3.48E5

49

17	paroxysmal-supraventric-tachycardia	20	6	6	4	2.83E5
23	noncardiogenic-pulmonary-edema	20	8.5	8	2	1.81E5
68	high-k	30	8	7	4	1.47E5
31	bronchospasm	30	8	7	4	1.47E5
62	low-mg	60	7	3	7	8.57E4
45	intrinsic-pathway-defects	60	7	3	5	4.37E4

Table 6.11. Selected Contingencies for k_p = 0.5 (30 seconds) for SPARC10 (k_t = 1.02)

The reactive planner model for it (kindly specified by my colleagu its designer, David Ash) states that the reactive plan built tends to implicit hierarchy with about 3 levels, with a roughly constant bra factor throughout. Actually distinguishing a child node in the im hierarchy is accomplished in the real hierarchy with a decision lis structure. According to this model, reaching a contingency in the plar for n contingencies takes roughly a constant time, equal to 3* time amount of time for a single test (assuming the tests take approxin constant time). This assumption can be made in our domain and for ou since tests which take much longer (e.g. laboratory tests) are to be inclu the main plan by the planner, to be performed regularly so that their always meaningful. This is generally the way physicians operate in rea settings. Therefore, for the purpose of our model, we can assume th length of a test is roughly given by the time a human operator needs i to retrieve and check a piece of data and to input it into the comput approximately 30 seconds. The reactive planner model also allows for a set of contingencies (say, three) to be hooked directly to the top o hierarchy, and thus to be reached by tests independently of the contingencies to be solved by this reactive plan. This is useful when the a few very time critical contingencies, and the rest are with a much s time pressure.

#	Contingency (Response would be the typical response for this event)	Resp. time	Consequences	Side-eff.	Likelihood	Criticality
34	et-tube-disconnection	2	10	2	4	4.2E12
18	ventricular-tachycardia	1	9	7	2	2.2E12

51

13
ventricular-fibrillation
1
10
8
1
6.1E11

35
kinked-et-tube
5
8
2
4
1.8E10

20
hypoxia
5
8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5

52

8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

Table 6.12. Selected Contingencies for k_p = 0.6 (36 seconds) for ExplorerII (k_t = 1.166)

The agent model only takes into account the slowdown of the syster to computational overhead. Simulations on two different platforms yielded significantly different results: if Guardian is run on Explo machines, the computational overhead is on average 16% for the sim time period we are interested in (approximately two hours of simulated on a SPARC10 workstation, this overhead is reduced to approximately 2%. 6.10 presents the results of running our entire framework, with the r planner and agent models described here, for the Guardian agent runr an ExplorerII platform. Table 6.11 presents the same results for a SPA workstation. We have run the same experiment for an estimated test ti 20% larger (36 seconds) and the results are presented in tables 6.12 ar for ExplorerII and SPARC10 respectively. In the second case, the system able to include about 75% more contingencies in the reactive plan. Als that in all cases the system was able to include about 66% more contin in the reactive plan to be run on the SPARC10.

#
Contingency (Response would be the typical response for this event)
Resp.
time
Conse-
quences
Side-
eff.
Likeli-
hood
Criti-
cality

53

34
et-tube-disconnection
2
10
2
4
4.2E12

18
ventricular-tachycardia
1
9
7
2
2.2E12

13
ventricular-fibrillation
1
10
8
1
6.1E11

35
kinked-et-tube
5
8
2
4
1.8E10

20
hypoxia
5
8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia

54

5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

1
myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism
10
8.5
7.5

55

3
2.13E7

6
hypovolemia
20
7
3
7
2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax
10
8
7
3
2.01E7

40
acute-hemolytic-transfusion-react
10
8.5
5
1
1.28E7

26
hemothorax
10
7
7
4
1.05E7

9
right-heart-failure
10
8
7
2
8.94E6

Table 6.13. Selected Contingencies for k_p = 0.6 (36 seconds) for SPARC10 (k_t = 1.02)

The sets of selected contingencies include the first as many as po contingencies in the order of their criticality value (table A4.2). They include the fourth contingency in table A4.2 because of the special tre of highly time pressured contingencies in the reactive planner mod allows for three contingencies to be reacted to faster than the rest - oth the set of contingencies might have included only the first contingencies, but very few others if any). Due to the decision tree form reactive plan, all leaves are reached in approximately the same time (3.4.1), so the set of contingencies selected is limited by the response time most time pressured contingency included (in our case 5 minutes, sin one and two minute contingencies are treated separately).

These experiments reinforce a few statements we have made along thesis. They show that the framework proposed here is useful in pruni set of contingencies for which the agent should prepare to react. Tl however necessary only in such domains where the number of contin is large enough to pose problems due to agent resource limitations (a have characterized such domains in chapter 2); Guardian and its dom typical in this respect. The performance of the enhanced agent improve the performance of the same agent without the benefit of our fram because in the latter case, the reactive planner would have prepa reactive plan to include all 68 contingencies, and due to its size, the r requirements for such a plan could not achieve reactions to the mos pressured contingencies in this set. The set of contingencies selected d on the characteristics of the agent and of its reactive planner (as rep by the agent model and reactive planner model). The more accurate models are, the better will be the use of agent resources made by the contingencies selected. Also note that the agent may exhibit different r behaviors, as defined by the reactive behavior model.

Our experiments also show that the necessary data for our framew be applicable exists in practice and can be acquired from experts in rea domains. The more difficult part of the knowledge acquisition process w identification of the set of contingencies possible in a given situation acquisition of the characteristic values for them was much easier, esp since their absolute values are less important than their relative order, the robustness of the framework).

The experiments described in this chapter and performed in dif domains requiring quite different types of human expertise (mundane highly skilled domains, etc.) demonstrate the applicability of our fran in the general types of domains described in chapter 2.

Chapter 7
Conclusions

Most research projects have their roots in one or two basic ques attempt (more or less successfully) to provide answers to these questior during this process usually generate many more new questions than a This thesis was no exception. In the next section, we present a summary answers which our work provides, and in the following section we enu a few questions raised and research avenues opened during our efforts solutions to the original problems stated in chapter 2.

7.1. Summary

Executing plans in the real world has long ago been recognized difficult and uncertainty-filled problem, due to contingencies generat interactions between the executing agent and its environment. Cond planning, reaction and dynamic replanning are all possible control mc solve this problem, but none of them alone is entirely suitable for agen limited resources working in complex environments. Therefore, the arises for a mechanism to select, from the set of possible contingencies domain, the subsets which should be treated using each of the prev mentioned control modes. In this thesis we have defined a framework to the subset of contingencies which are best suited for reactive respons framework's decisions are based on the plan situation under considerati characteristics of the contingencies and of an expert model specifying as well as on the reactive planner and agent models. A behavior determines the type of reactive behavior to be exhibited by the ager these models are designed as application-dependent plug-in modules attached to our framework, thus substantially increasing its generalit applicability across domains and types of agents. The decision of whet prepare a reaction to a given contingency or not is taken while cons the entire set of contingencies that may appear in that situatio relationship with the limitations of the agent's execution time resourc have justified a few theoretical claims about our framework (includin optimal use of agent's resources), and then we have verified experimentally. We have also demonstrated other properties of the fran the most important being that the reactive behavior of an agent usi framework has the agreement of the experts in the field.

A couple of extensions to our framework were also discussed. The one involves a similar framework to decide on the subset of contingenc which to prepare a conditional branch (all the way to the final goal) plan. The second involves a proposal for a knowledge represent formalism for the types of knowledge involved in our framew contingencies, reactions and situations. It was designed to facilitate structuring and manipulation of this knowledge, as well as to facilitate t of automatic knowledge acquisition and learning techniques to cope wi explosion of the related knowledge in complex domains. However, both extensions were discussed only at a theoretical level and, as stated in th section, they need more work in order to be fully understood and fo: potential to be fully used.

7.2. Future Work

It is unfortunate (or maybe actually very fortunate) that a thesis encompass an entire research career. Unfortunate because while tryi solve the originally stated problems, there are so many new problem arise and which I would have liked to address. Fortunate because I am su while trying to address these new issues, many other problems would and then no thesis would ever be finished. We shall briefly overview i: section a few of the research issues which came up while solving the pr mentioned before.

Two already stated issues are the extensions to our framework above. The first involves the framework for deciding whether to prepare conditional branch for a contingency. While we have defined the ge framework in section 3.5, there are many details that still have to be sor before a usable framework like the one for reactions can be obtained function computing the conditional planning value of a contingency m identified and tested, and the values for its parameters must be specifie normal behavior model (and possibly for other types of behavior m Guidelines for specifying the planner model and especially the agent (from the perspective of conventional plan execution) must also be set.

The second issue involves the knowledge representation forma proposed in chapter 4. Since specifying the nonterminals of the gra imposes some additional burden on the experts, it would be very help devise a set of knowledge acquisition and learning tools to help the ex this task. We believe that the best results here can be achieved by com automatic learning methods with interactive knowledge acquisition (similarly with the methods used in [Dabija & al., 1992a]). Such an ap would better use the potential for bias shifting [Utgoff, 1988] and co classification that this knowledge representation formalism is appropria

Another open research issue related to our framework is its pot integration with case based reasoning and planning techniques. Figu presents the possible information flow in such a system. The ag knowledge base (contingencies and associated reactions in specific situ may be organized as a library of cases. The agent may also have a libr reactive plans already built (each reactive plan built, may be cached in library), organized by the situations in which they may apply. New kno may be added at any time to the case library, and each time an a encountered situation arises, the reactive plan that may already exist plan library is combined with any new contingency-reaction pairs app in that situation that have been included in the agent's knowledge bas the last use of this reactive plan. Our framework will decide, for each situation, which are the best contingencies for which reactions shou included in the updated reactive plan. If no new relevant knowledge applicable in the current situation) has been added to the knowledge since its last use, this reactive plan may be used without any changes or

Many issues arise here related to the independent management of th libraries (knowledge structuring, and "forgetting") as well as relationships between them. There are also interesting research issues to the problem of acquiring the knowledge for the two libraries: knov for each of them may be acquired from an expert (and here inter knowledge acquisition techniques may be used) or from the agent's domain experience.

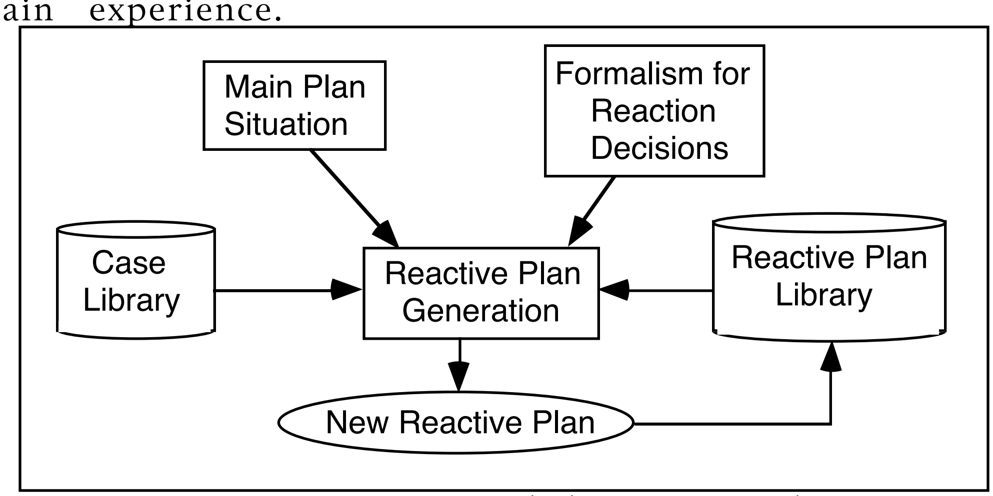

Figure 7.1. Extended system architecture

In domains where strong theories about possible contingencies these theories can be used to anticipate all the contingencies that may for situations along the plan, and to specify their characteristics. Howe most domains with which we are concerned, such theories either do no or they are very weak (e.g. cover the domain only partially, or can ant only certain kinds of events all over the domain). In such cases, the age generate prototype cases (akin to the cases in the case library) and p solutions for them. They may then be evaluated and comparec corresponding actual cases, and the differences may be used to impro weak domain theory that has generated them in the first place.

In this thesis we have also introduced a formalism to describe rea behavior models. As we have shown in chapter 6, most of the human r behavior models described in the literature can be conveniently expres our framework, which therefore provides a possible vehicle for the ex of information on this subject. However, we have only touched the tip iceberg in this respect. Considerably more research is needed to refin formalism so that it can be really useful for providing comp characterizations of these behavior models and therefore become use attempts to correct or influence human behaviors in critical domain nuclear power plant supervision or aircraft flying. For example, in or better model the differences between behaviors like social responsibili individualism, the consequences dimension of the criticality space m split into two components: (i) internal-consequences (which directly our agent) and (ii) external-consequences (effects of not responding contingency, over other agents in the environment).

As stated at the beginning of this section, the range of open pro suggested by this research is very wide, and we believe that at least p them are worth further investigation.

Appendix 1
System Architecture

We briefly present here the way our framework is to be integrate the general architecture of an agent with planning, reaction and mor capabilities. We assumed a modular system, in which each component c principle, be plugged in and out and the agent's performance should gracefully. For example, if the agent is to operate without a reactive p then it will be able to respond only to the contingencies for which con branches have been prepared by the planner, while if it is to operat with a reactive planner, then the agent should be able to react to a contingencies for which it has reactions prepared for, but may never the overall goal since it lacks the main plan to do it. The framework to whether to prepare to react may be regarded as another such module, when present, ensures that the agent is better prepared to cope wi different contingencies that may appear during its plan execution.

An alternative view is that the other agent modules (the plar reactive planner, execution mechanisms, knowledge base, the expert and the behavior model) are all independent modules which can be p into, and out of, the framework discussed in the thesis. The framewor defined in a general manner such that all these modules are parameter will change the outcome of the analysis, but the general principles pr in chapters 3 and 4 and the theoretical analysis in chapter 5 remain a (since they all were done independent of any particular such module).

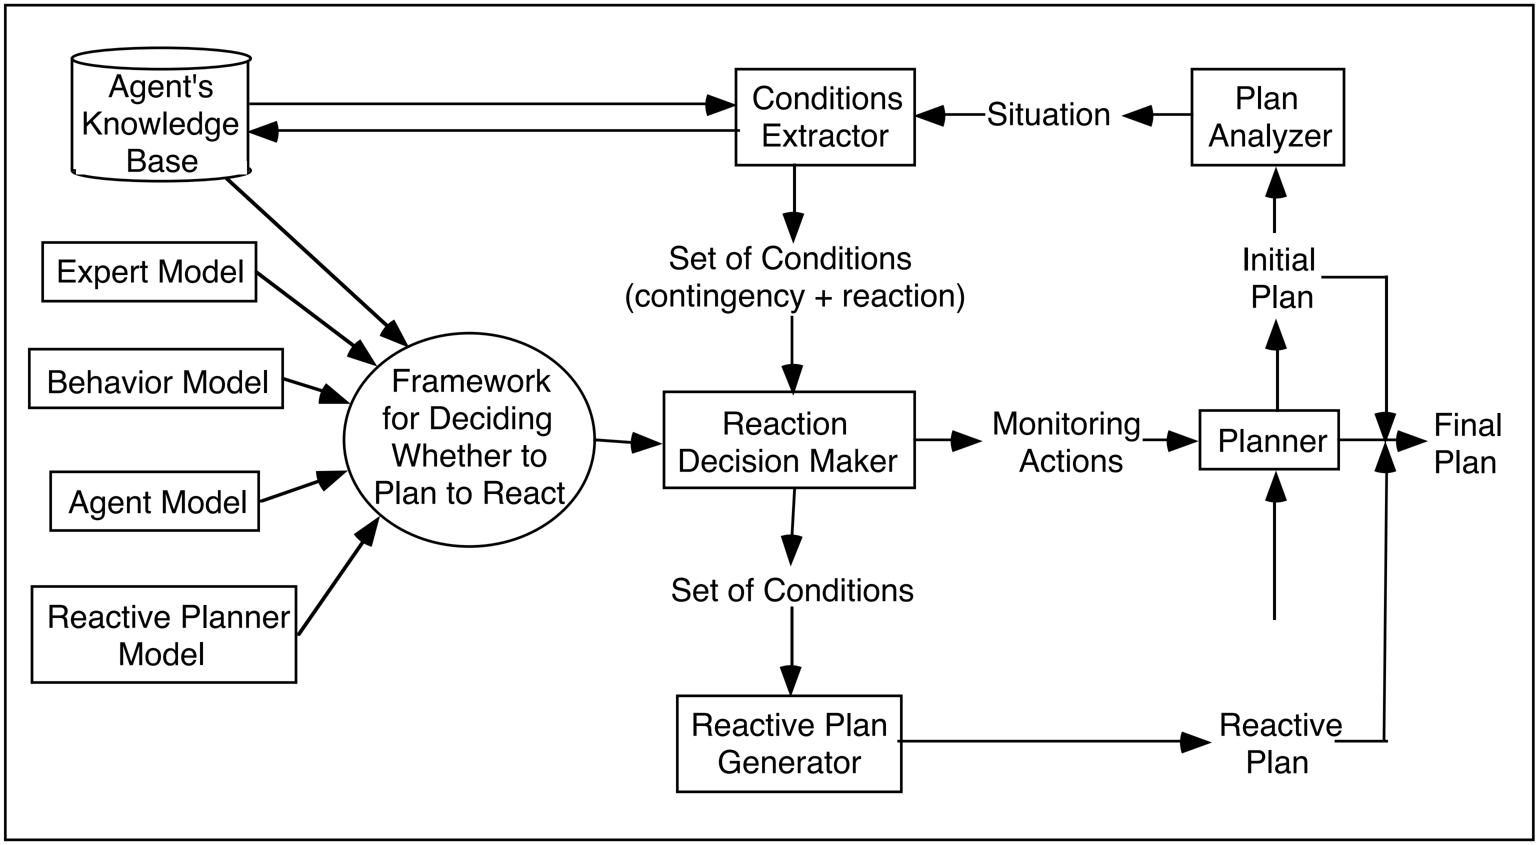

Figure A1.1. System Architecture and Information Flow

Figure A1 presents how all these modules fit together in a "comp agent, as well as the information flow during the plan modification p We assume this process starts when the planner has produced a complete (conditional) plan to solve a given problem. In order to identify the sit that may generate contingencies in the plan, the plan analyzer scans t and for each stage (situation) searches the agent's knowledge base for the set of contingencies that may appear in that situation, and their appr reactions. Each situation for which there are known contingencies wi further analyzed to prepare reactive plans for it.

All relevant contingencies found in the agent's knowledge base by the contingency extractor for a certain situation are passed on to the reaction decision maker which uses our framework presented in chapter 3, toget with an expert model, a behavior model, the agent model (corresponding to the execution capabilities of this agent), and the reactive planner m (corresponding to the reactive planner available to this agent), to selec contingencies for which reactive responses should be prepared by reactive plan generator. The reactive plan is passed back to the planne together with monitoring actions to be included in the plan. The reacti is eventually attached to the context-specific plan and the next stage plan will be subsequently analyzed.

This entire process is performed first at planning time, before the starts executing the main plan, and is repeated each time the agent is fc dynamically replan its actions (and generate a new main plan) durir execution phase because of a major failure in executing the initial main

One agent with such an architecture with which we have condu demonstrations of our framework is the Guardian agent (for moni patients in an intensive care unit) [Hayes-Roth, 1990]. The results of demonstrations are discussed in section 6.4.

Appendix 2
Knowledge Representation in the Car-Driving Domain

We continue here the example started in section 4.2, with hierarchical vocabularies and the corresponding grammars for repre reactions and situations in the car driving domain. While we do not p specify the complete vocabularies for this domain, the ones that are here are sufficient to represent all the examples encountered in chapte well as the experiments discussed in chapter 6 for the driving domain are also enough to represent a good deal more knowledge from this dom

Figure A2.1 presents the hierarchical vocabulary for represen reactions in the car driving domain. This hierarchy is equivalent (acc to the formalism discussed in chapter 4) to the following grammar:

```
G = (N, T, P, S), where:
    N = { Reaction, Brake, Steer, Other, Left, Right, Hard, Gently,
          Adjust_Radio }
    T = { B.Hard, B.Gently, B.None, Left&Hard, Right&Hard, Left&Gently,
          Right&Gently, None, Turn_on_Lights, Adjust_Volume,
          Adjust_Station, Open_Window }
    P = { Reaction -> Brake - Steer | Other
          Brake -> B.Hard | B.Gently | B.None
          Steer -> Left | Right | Hard | Gently | None
          Left -> Left&Hard | Right&Hard
          Right -> Right&Hard | Right&Gently
          Hard -> Left&Hard | Right&Hard
          Gently -> Left&Gently | Right&Gently
          Other -> Turn_on_Lights | Adjust_Radio | Open_Window | ...
          Adjust_Radio -> Adjust_Volume | Adjust_Station }
    S = Reaction
```

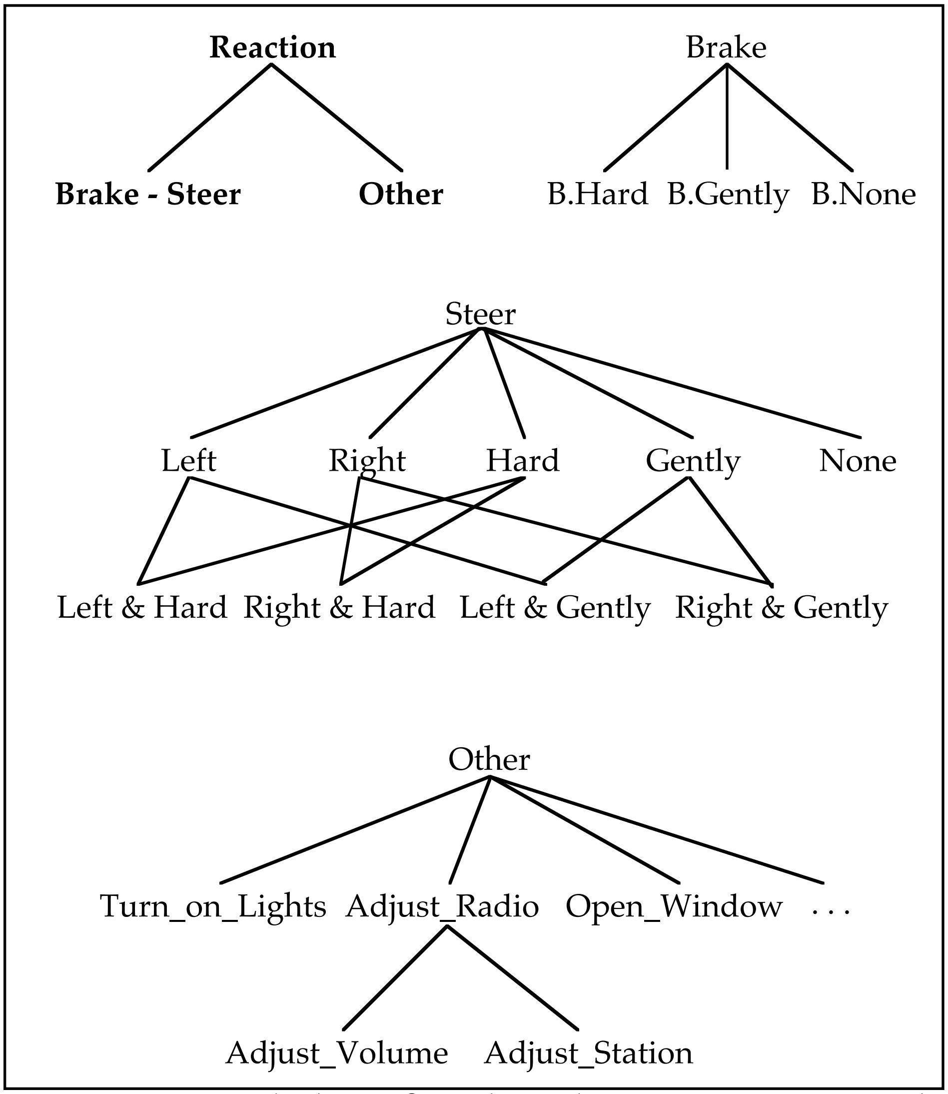

Figure A2.1. Vocabulary for describing reactions in the driving doma

Every reaction specified in table 3.1 can be obtained through a nu of different derivations in this very small and simple grammar. Also, other reactions in the driving domain can be expressed using this voc (this is generally true especially for reactions, since there are usually a set of actions in a domain which can make up useful reactive plans in domain). Since general reactions are often enough to be specified, derivations may be stopped at those levels where the reaction expressed sentential form obtained thus far "contains" (according to the order r defined in chapter 4) all the elementary reactions acceptable in that si For example, if the agent only needs to reduce speed somewhat, than " may be sufficient, without qualifying the action further.

Here is an example of deriving the reaction "Brake hard and s right" to the first contingency in table 3.1 ("Child runs from right 2 front of car"):

Reaction ->

Brake - Steer ->

B.Hard - Steer ->
B.Hard - Right.

This derivation has already been stopped before reaching a sent form made up only of terminals in the vocabulary, since the "R nonterminal could have been further refined to one of the two ter given by the production:

Right -> Right&Hard | Right&Gently.

It therefore represents a set of possible reactions, contained in description (i.e. derivable from it).

Figure A2.2 presents the hierarchical vocabulary for represen situations in the car driving domain.

Some productions (both shown in figure A2.2 and omitted) may realized through identification functions, as shown in chapter 4. For e the grammar symbols Slow, Medium, Fast, can be considered nonterm (instead of terminals like in this example), and the actual values of the can be considered terminals.

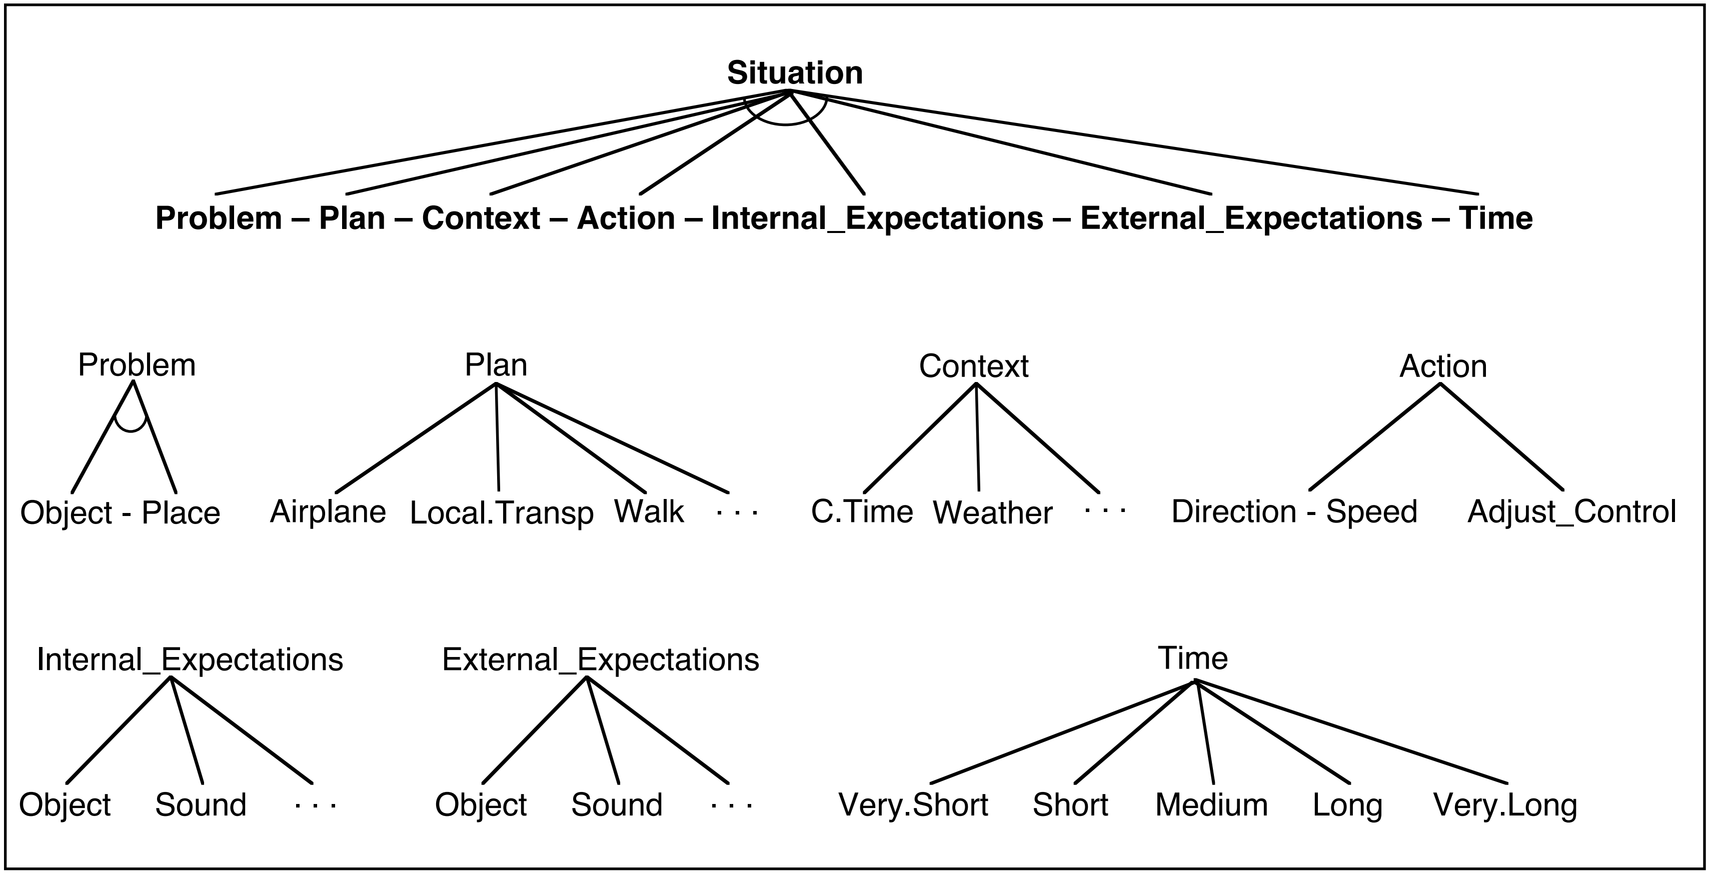

Figure A2.2a. Vocabulay for describing situations in the driving domain

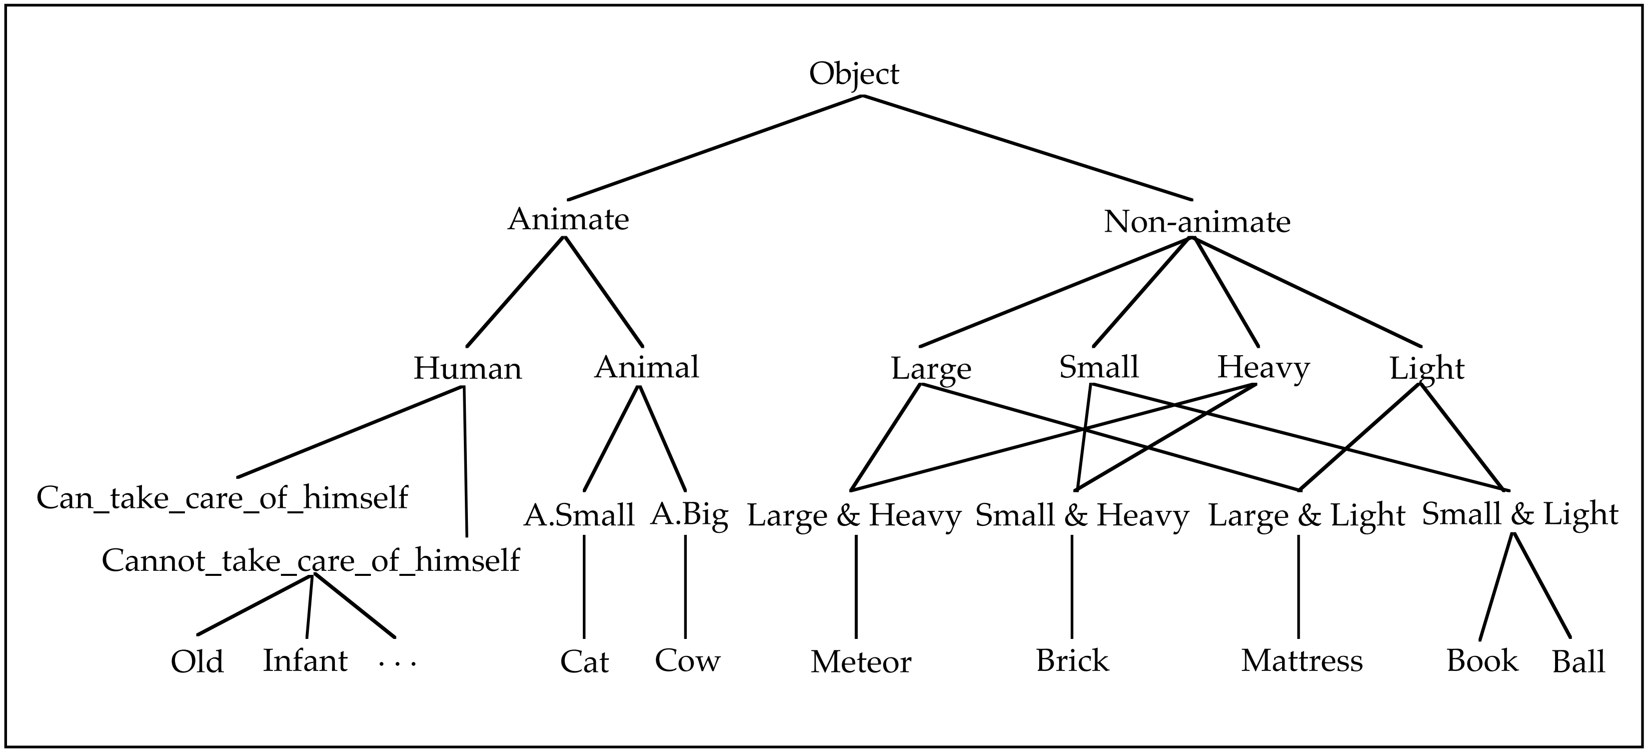

Figure A2.2b. Vocabulay for describing situations in the driving domain (continued)

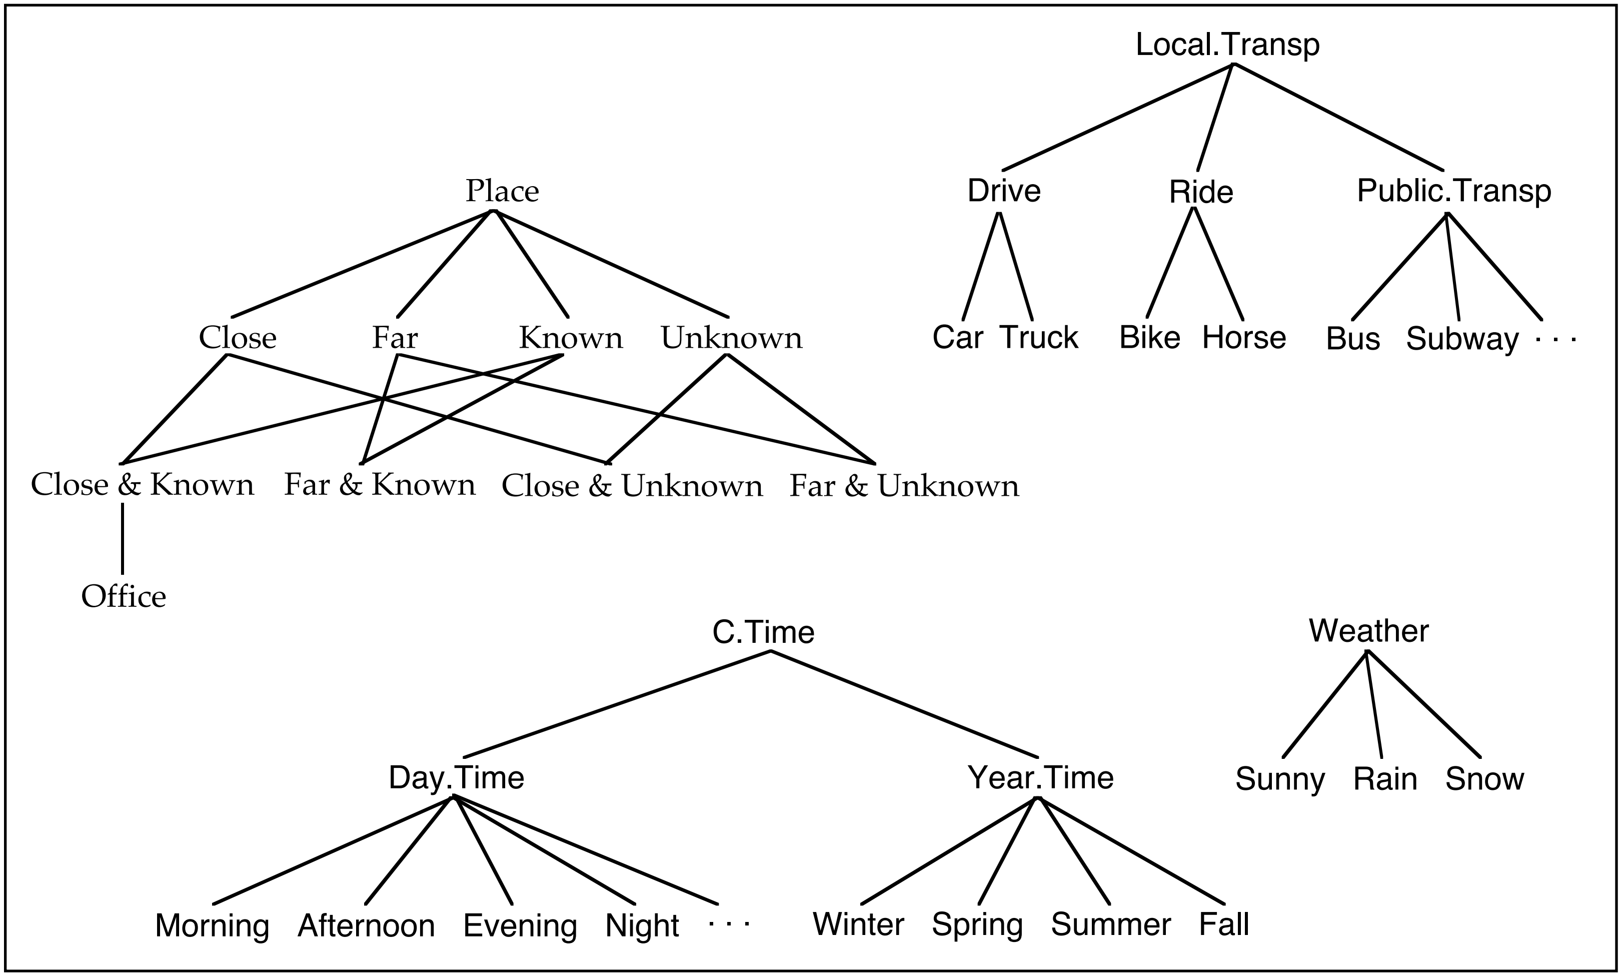

Figure A2.2c. Vocabulay for describing situations in the driving domain (continued)

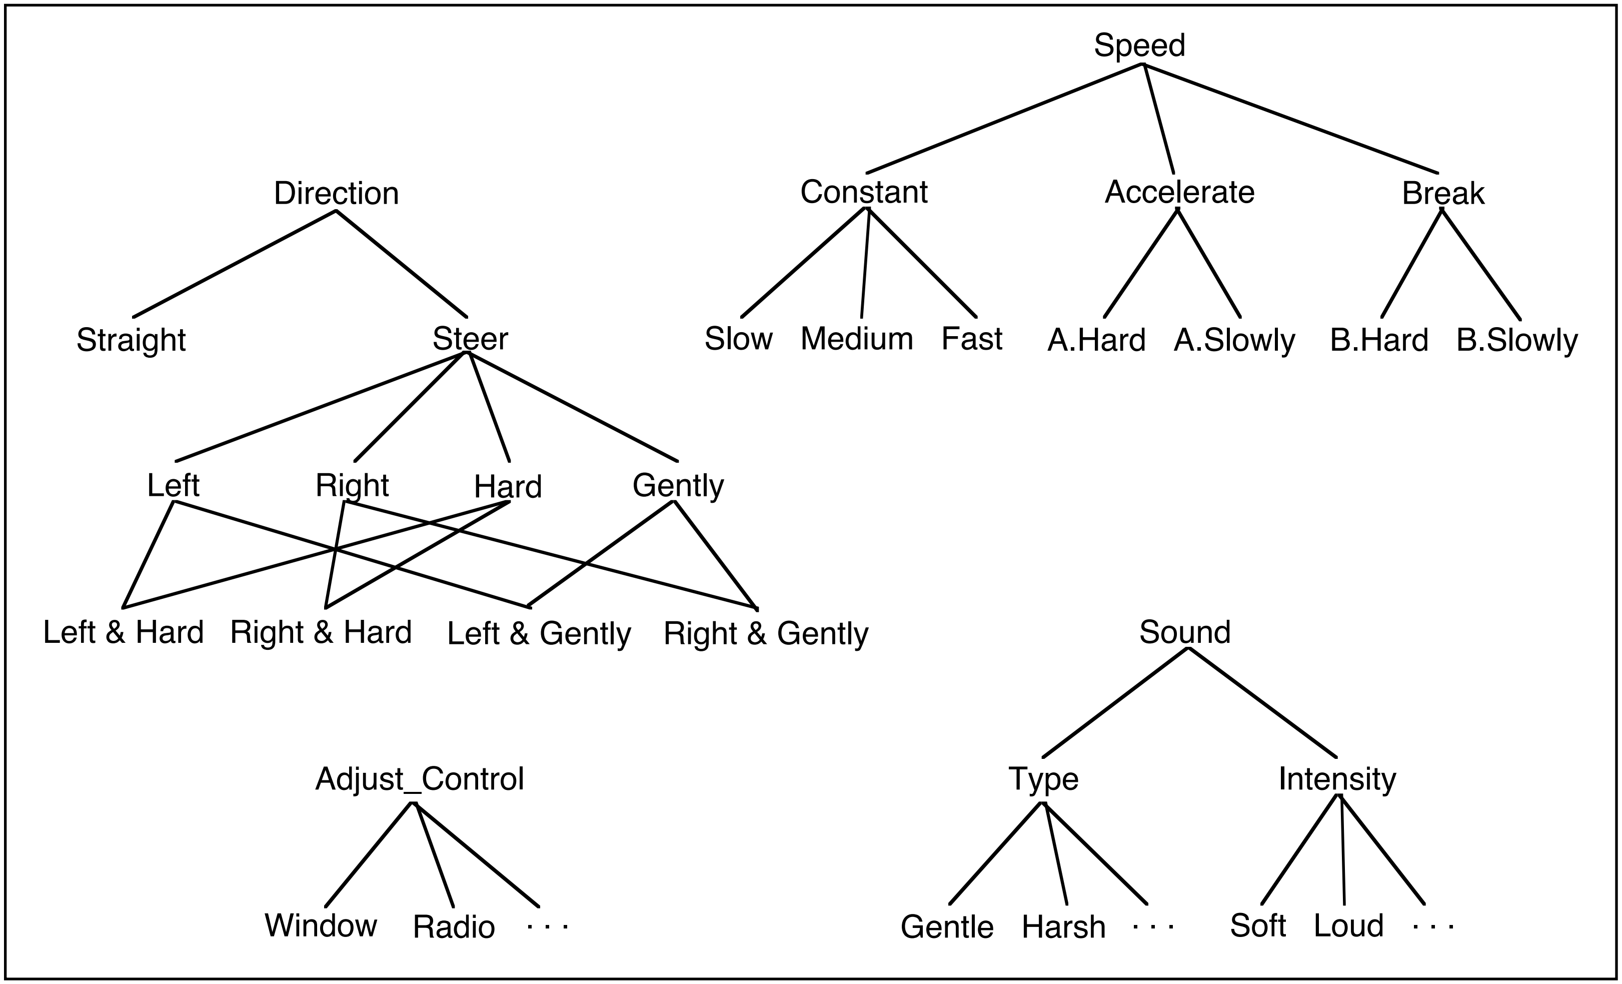

Figure A2.2d. Vocabulay for describing situations in the driving domain (continued)

An example of such a function may then be:

$$Slow = f_S\,(speed) = 5\ mph < speed < 20\ mph,$$

which can be used to perform the transition over the edge linking "Slow the actual terminal, say "speed = 15 mph".

We have collapsed the seven vocabularies for representing values the seven dimensions of the situation space into a single vocabulary, w help of the first production of the grammar. Alternatively, we could specified seven independent grammars, by throwing out the first pro and the nonterminal Situation; each of these grammars would have h starting symbols the nonterminals: Situation, Problem, Plan, Context, A Internal_Expectations, External_Expectations, Time (respectively), productions all the productions which can be reached from their res start symbols using the productions of the reunited grammar, an nonterminals and terminals all those from the large grammar whic involved in the productions of each respective grammar.

The hierarchy in figure A2.2 is equivalent (according to the form discussed in chapter 4) to the following grammar:

G = (N, T, P, S), where:

N = { Situation, Problem, Plan, Context, Action, Internal_Expectations
External_Expectations, Time,
Object, Animate, Human, Cannot_take_care_of_himself, Anima
A.Small, A.Big, Non-animate, Large, Small, Heavy, Light,
Large&Heavy, Small&Heavy, Large&Light, Small&Light,
Place, Close, Far, Known, Unknown, Close&Unknown,
Local.Transp, Drive, Ride, Public.Transp,
C.Time, Day.Time, Year.Time, Weather,
Direction, Steer, Left, Right, Hard, Gently,
Speed, Constant, Accelerate, Break,
Adjust_Control,
Sound, Type, Intensity }

T = { Airplane, Walk,
Very.Short, Short, Medium, Long, Very.Long,
Can_take_care_of_himself, Old, Infant, Cat, Cow, Meteor, Brick,
Mattress, Book, Ball,
Office, Far&Known, Close&Unknown, Far&Unknown,
Car, Truck, Bike, Horse, Bus, Subway,
Morning, Afternoon, Evening, Night,
Winter, Spring, Summer, Fall,
Sunny, Rain, Snow,
Straight, Left&Hard, Right&Hard, Left&Gently, Right&Gently,
Slow, Medium, Fast, A.Hard, A.Slowly, B.Hard, B.Slowly
Window, Radio,
Gentle, Harsh, Soft, Loud,
. . . }

P = { Situation -> Problem - Plan - Context - Action -
Internal_Expectations - External_Expectations - Time
Problem -> Object - Place
Plan -> Airplane | Local.Transp | Walk | . . .
Context -> C.Time | Weather | . . .
Action -> Direction - Speed | Adjust_Control
Internal_Expectations -> Object | Sound | . . .

External_Expectations -> Object | Sound | . . .
Time -> Very.Short | Short | Medium | Long | Very.Long
Object -> Animate | Non-animate
Animate -> Human | Animal
Human -> Can_take_care_of_himself |
    Cannot_take_care_of_himself
Cannot_take_care_of_himself -> Old | Infant | . . .
Animal -> A.Small | A.Big
A.Small -> Cat | . . .
A.Big -> Cow | . . .
Non-animate -> Large | Small | Heavy | Light
Large -> Large&Heavy | Large&Light
Small -> Small&Heavy | Small&Light
Heavy -> Large&Heavy | Small&Heavy
Light -> Large&Light | Small&Light
Large&Heavy -> Meteor | . . .
Small&Heavy -> Brick | . . .
Large&Light -> Mattress | . . .
Small&Light -> Book | Ball | . . .
Place -> Close | Far | Known | Unknown
Close -> Close&Known | Close&Unknown
Far -> Far&Known | Far&Unknown
Known -> Close&Known | Far&Known
Unknown -> Close&Unknown | Far&Unknown
Close&Unknown -> Office | . . .
Local.Transp -> Drive | Ride | Public.Transp
Drive -> Car | Truck
Ride -> Bike | Horse
Public.Transp -> Bus | Subway | . . .
C.Time -> Day.Time | Year.Time
Day.Time -> Morning | Afternoon | Evening | Night | . . .
Year.Time -> Winter | Spring | Summer | Fall
Weather -> Sunny | Rain | Snow
Direction -> Straight | Steer
Steer -> Left | Right | Hard | Gently
Left -> Left&Hard | Left&Gently
Right -> Right&Hard | Right&Gently
Hard -> Left&Hard | Right&Hard
Gently -> Left&Gently | Right&Gently
Speed -> Constant | Accelerate | Break
Constant -> Slow | Medium | Fast
Accelerate -> A.Hard | A.Slowly
Break -> B.Hard | B.Slowly
Adjust_Control -> Window | Radio | . . .
Sound -> Type | Intensity
Type -> Gentle | Harsh | . . .
Intensity -> Soft | Loud | . . . }

S = Situation

Most of the driving domain situations encountered during this t can now be obtained through a number of different derivations ir grammar. Also, many other situations in the driving domain can be ex using this vocabulary. Clearly, this vocabulary is not enough to descri possible contingencies in the driving domain. It was not our goal to p such a vocabulary and grammar. However, it can be easily extende encompass, in the same domain, any other desired situation which can represented yet.

Contingencies and reactions are, in general, associated with set situations. Therefore, general situations are most often enough to be sp and the derivations may be stopped at those levels where the sit expressed by the sentential form obtained thus far "contains" (accord the order relation defined in chapter 4) all the elementary situations to the contingency or reaction apply. This knowledge structuring prope the representation formalism is most important here, since it helps the explosion of the situations in the domain, ensuring the representab the knowledge needed for our planning-to-react decision framework in domains.

While most situations encountered in chapter 3 can be derived in formalism, it also supports the derivation of many other situations f driving domain. In fact, just by enlarging the set of terminals, the num situations expressible with this small grammar becomes very large in This fact underlines the most important advantage of this represer formalism, namely imposing a (hierarchical) structure on the set of p situations in the domain, which then makes them much easier to be managed, analyzed and reasoned about.

Appendix 3
Anesthesiology Domain Experiments

In order to demonstrate the applicability and scalability of the re decision framework presented in chapter 3, we have run demonstrati one other domain than those described in chapter 6. We briefly descril these demonstrations. The domain is anesthesiology, and I am indebted David Gaba for letting me benefit from his time and knowledge by servi role of the domain expert both for the knowledge acquisition task, as v for the evaluation phase of the experiments. Working in a professional of high expertise, we have used this time a single expert to provide u necessary knowledge (in contrast with the driving domain where we acquired it through a statistical analysis of the opinions of a group of in the domain, as explained in section 6.1).

Table A3.1 lists the set of 13 contingencies selected for this experi together with the reactions for each of them (in the "random" order s by the expert), for the following situation:

Problem: Anesthetize patient for bowel obstruction
Plan: Induce anesthesia [rapid sequence induction]
Context: Middle of the night, emergency case, patient has coron artery disease (moderate) and chronic obstructiv pulmonary disease (severe)
Ext. Expect.: Change in vital signs
Int. Expect.: Patient becomes unresponsive to commands
Action: Rapid sequence induction (Pentothal and Succinylcholi have just been administered)
Time$_S$: 60 seconds.

The expert was asked to translate his qualitative feelings i quantitative values, and to concentrate more on relative values than absolute values he was going to specify. The expert was not asked to orc contingencies as he feels would be appropriate for a normal behavior. we have presented him with the system's results and ask him to evalu behavior recommended by our framework. The knowledge acquired fro expert was for the following contingency characteristics: time to re (real values in seconds), criticality, side-effects, and likelihood (all these on a scale of [0,10]).

	Contingency	Reaction
1	Patient vomits	

2

Turn head; suction mouth; intubate

2
Patient does not "fall asleep"
Check IV and syringe; give more drug

3
Muscle fasciculations (twitching 2" to drug)
Ensure patient does not fall asleep

4
Decreased blood pressure
Increase IV rate; administer vasopressor

5
Increased heart rate
Consider deeper anesthesia or β blocker

6
Cardiac Arrest
ACLS (Advanced Cardiac Life Support)

7
Meteor strikes OR
Move patient out of OR

8
Failure of pipeline oxygen supply
Switch tanks ON; disconnect pipeline

9
Failure of 1° and backup electric power
Obtain flashlight

10
Inability to intubate trachea
Ventilate by mask if possible; emer-gency procedures for difficult airwa

11
Message from PACU about previous patient
Listen to the message

12
Severe bronchospasm (wheezing)
Ensure correct intubation; treat with bronchodilators

13
O2 saturation decreases to < 90%
Ventilate by mask or tube with 100% O2

Table A3.1. Contingencies for the anesthesia domain experiments

We have also asked the expert to calibrate his data by supplying v
for the expert model parameters for the recommended behavior model.

values were: 1 second for minimum real time (corresponding to T_{max} minutes for maximum real time threshold (corresponding to T_{min}), minimum likelihood (L_{min}), and 2.3 for the difference between conseq and side-effects (CS_{min}). We did not ask the expert to actually give function to translate from real time to time pressure, but rather we specified it ourselves, in such a way as to include most of the time press the interval [0,10]. The function we came up with is:

$$f_{tc} = k \ / \ time_{rc} = 50 \ / \ time_{rc} \ .$$

We have experimented with significantly different values for k (bet [10,100]) and the results obtained were remarkably similar (actually mo identical) with the ones reported here. However, we have settled for the k = 50, for the reason stated above (all but one time pressure value between [0,10], with a reasonable spread in this interval). The results knowledge acquisition process in this domain are summarized in table A the same order as the previous table).

Contingency		$time_{rc}$	$time_p$	consequence	side-effect	likelihood
1	vomit	15.0	3.33	8.0	2.0	7.0
2	not fall asleep	45.0	1.11	7.0	4.0	4.0
3	muscle fascic.	100.0	0.5	3.0	1.0	8.0
4	decreased BP					

4

15.0
3.33
8.0
5.0
6.0

5
increased HR
15.0
3.33
6.0
6.0
7.0

6
cardiac arrest
5.0
10.0
10.0
7.0
2.0

7
meteor
0.1
500.0
9.0
7.0
0.01

8
O2 supply fails
30.0
1.67
8.5
5.0
1.0

9
power failure
30.0
1.67
6.0
5.0
1.0

10
can't intubate
10.0
5.0
9.5

5

8.0
5.0

11
PACU message
200.0
0.25
1.0
1.0
4.0

12
bronchospasm
25.0
2.0
9.0
7.0
6.0

13
O2 sat < 90%
15.0
3.33
8.0
4.0
6.0

Table A3.2. Data values for the anesthesiology domain experiments

We have first run the "normal" behavior model on these continge
The values for the criticality function parameters given by the bel
model were the same as for the driving domain:

$$p_1 = 5, p_2 = 1, p_3 = 0, p_4 = 0, p_5 = 3, p_6 = 2,$$

with the parameters specified by the expert model (and discussed above
to those given in section 6.2. Table A3.3 summarizes the results of this ru
contingencies are this time numbered in the order specified by the cr
function for this case, which we shall call from now on t
"system-recommended" order (since it was obtained by running the
with the recommended behavior model). There are two possible mon
thresholds, since there are two significant gaps in the sequence of
returned by the criticality function.

Contingency

Criticality
Monitor

1
cardiac arrest
5.95E8

6

yes

2
vomit
9.22E7
yes

3
can't intubate
4.07E7
yes

4
O2 sat < 90%
2.96E7
yes

5
decreased BP
1.76E7
yes

6
increased HR
1.47E6
yes

7
bronchospasm
8.24E5
yes

8
not fall asleep
2.82E4
??

9
O2 supply fail
2.13E4
??

7			
10	power failure	2.77E3	??
11	muscle fascic.	4.77E2	??
12	PACU messg	0.19	
13	meteor	0.00	

Table A3.3. Criticality values for the "normal" behavior model, for the anesthesiology domain experiments

As mentioned before, the expert was not required to order contingencies by reaction value according to his belief of what recommended behavior should be like. However, when presented wit results, he characterized them as "definitely reasonable". This sho significant portability of the behavior model and of all the parameter for the criticality function, across domains (since the driving anesthesiology domains are significantly different in nature, and the have specified their knowledge in the two domains independent of other).

We have then run our framework on this data, for all the o behavior models defined in section 6.3. We summarize in table A3.4 the we have used for the criticality function parameters in each run fo domain. Note that all the behavior model parameters (p_1 to p_6) have received identical values for the two domains. Also most of the expert model par are unchanged, and the changes reflect the different calibrations experts when they have specified the data.

	Behavior Model						Expert Model
Behavior	p_1	p_2	p_3	p_4	p_5	p_6	

T_{max} T_{min} CS_{max} L_{min}

Recommended
5 1 0 0 3 2
50.0 0.028 2.3 1.0

Antiauthority
5 1 0 0 3 0

Impulsivity
0 0 0 0 0 3
5.0 5.0

Invulnerability
5 1 0 0 3 2
5.2

Macho
4 1 0 0 0 3
10.0

Resignation
5 1 0 0 3 2
2.0

Risk-averse
2 2 4 2 1 1

Liability conscious
3 3 1 2 1 2
500.0 0.0 0.0

Social responsibility
4 3 0 0 4 3

Table A3.4 Representing Behavior Models

Table A3.5 summarizes the results of these experiments. We have shown the reaction values produced by the criticality function. Their a values have no meaning whatsoever; what matters are their relative

(and only within the same behavior model), which represent the r
value of reacting to one contingency vs. another in a same situation. Fo
behavior, monitoring thresholds were set (for the expert model) in regi
the contingency space where there are big gaps among the reaction val
the contingencies ordered by criticality. The thresholds are represent
thicker lines separating the contingencies for each behavior into two o
sets (in many cases, two possible places were indicated for this threshol

Behavior Model 1

Behavior Model 2

Behavior Model 3

(Recommended)

(Antiauthority)

(Impulsivity)

1
cardiac arrest
5.95E8
1
cardiac arrest
1.48E8
3
can't intubate
1.25E2

2
vomit
9.22E7
2
vomit
1.88E6
11
muscle fascic.
22.62

3

10

can't intubate
4.07E7
3
can't intubate
1.62E6
2
vomit
18.52

4
O2 sat < 90%
2.96E7
4
O2 sat < 90%
8.23E5
6
increased HR
18.52

6
decreased BP
1.76E7
5
decreased BP
4.90E5
7
bronchospasm
14.69

5
increased HR
1.47E6
6
increased HR
3.00E4
5
decreased BP
14.69

7
bronchospasm
8.24E5
7
bronchospasm
2.28E4
4
O2 sat < 90%
14.69

8

11

not fall asleep
2.82E4
9
O2 supply fail
2.13E4
8
not fall asleep
8.00

9
O2 supply fail
2.13E4
10
power failure
2.77E3
12
PACU messg
8.00

10
power failure
2.77E3
8
not fall asleep
1.76E3
1
cardiac arrest
2.82

11
muscle fascic.
4.77E2
11
muscle fascic.
7.45
9
O2 supply fail
1.00

12
PACU messg
0.19
12
PACU messg
1.1E-2
10
power failure
1.00

13

12

meteor
0.00
13
meteor
0.00
13
meteor
0.00

Table A3.5 Reactive Behavior Experiments for Anesthesiology

Behavior Model 4

Behavior Model 5

Behavior Model 6

(Invulnerability)

(Macho)

(Resignation)

2
vomit
9.22E7
1
cardiac arrest
8.00E5
7
bronchospasm
8.24E5

4
O2 sat < 90%
2.96E7
3
can't intubate
7.42E5
8
not fall asleep

13

2.82E4

5
decreased BP
1.76E7
2
vomit
3.28E5
9
O2 supply fail
2.13E4

6
increased HR
1.47E6
6
increased HR
2.54E5
10
power failure
2.77E3

7
bronchospasm
8.24E5
5
decreased BP
2.13E5
11
muscle fascic.
4.77E2

1
cardiac arrest
2.44E4
4
O2 sat < 90%
2.13E5
12
PACU messg
0.19

3
can't intubate
6.38E3
7
bronchospasm
3.11E4
1
cardiac arrest

14

0.00

11
muscle fascic.
4.77E2
8
not fall asleep
6.82E2
3
can't intubate
0.00

8
not fall asleep
1.68E2
11
muscle fascic.
96.00
13
meteor
0.00

9
O2 supply fail
1.46E2
9
O2 supply fail
65.58
2
vomit
0.00

10
power failure
52.65
10
power failure
46.29
5
decreased BP
0.00

12
PACU messg
0.43
12
PACU messg
0.25
4
O2 sat < 90%

15

0.00

13
meteor
0.00
13
meteor
0.00
6
increased HR
0.00

Table A3.5 Reactive Behavior Experiments for Anesthesiology (continue

Behavior Model 7

Behavior Model 8

Behavior Model 9

(Risk-averse)

(Liability conscious)

(Social responsibility)

1
cardiac arrest
7.3E10
13
meteor
7.0E10
2
vomit
1.0E11

3
can't intubate
5.3E10
1

16

cardiac arrest
4.2E10
1
cardiac arrest
6.3E10

7
bronchospasm
5.13E9
3
can't intubate
2.4E10
4
O2 sat < 90%
2.1E10

5
decreased BP
2.38E9
5
decreased BP
3.05E9
3
can't intubate
1.3E10

6
increased HR
1.20E9
4
O2 sat < 90%
2.47E9
5
decreased BP
1.0E10

4
O2 sat < 90%
9.90E8
7
bronchospasm
1.61E9
7
bronchospasm
8.61E8

9
O2 supply fail
1.32E8
2

vomit
1.54E9
6
increased HR
2.55E8

2
vomit
6.61E7
6
increased HR
7.78E8
8
not fall asleep
2.64E7

8
not fall asleep
3.97E7
8
not fall asleep
1.93E7
9
O2 supply fail
5.36E6

10
power failure
2.49E7
9
O2 supply fail
1.50E7
10
power failure
1.97E5

11
muscle fascic.
1.23E3
10
power failure
1.99E6
11
muscle fascic.
2.95E5

12
PACU messg
2.30
11

18

muscle fascic.
1.48E4
12
PACU messg
6.99

13
meteor
0.00
12
PACU messg
2.30
13
meteor
0.00

Table A3.5 Reactive Behavior Experiments for Anesthesiology (continue

The numbering of contingencies for each behavior model in table is the same as for the recommended behavior. This was done in ord facilitate comparisons of each behavior model with the "normal" one.

In chapter 5, we have defined a behavior model to be an o relationship on the set of contingencies associated with a situation. The in these experiments, we only concentrate on the ordering of conting by reaction value (and sometimes relative values of the criticality fu but never on its absolute values), and ignore any issues related to the planner model and the agent model, that is we ignore the final decis applying the framework to a set of contingencies. This is consistent wi purpose of our demonstrations here, since any specific agent (with a reactive planner and resource limitations) may exhibit any of the re behaviors discussed, depending only on the order in which its behavior recommends the contingencies for consideration to be reacted to, and the actual components and resources of the agent.

The results of these demonstrations require a certain amoun interpretation (this is necessary especially since the definitions of behavior models are generally based on execution time types of rea while we attempt here to implement them at planning time). For examp the antiauthority behavior model, the order of contingencies does not much, since here almost all contingencies considered are covered regulations or procedures; only "not fall asleep" goes down since after a is precisely what we want to achieve and is therefore best covered procedures in this case. In the invulnerability case, "cardiac arrest" and "can' intubate" fall significantly (possibly even below the monitoring thre because they are not likely enough in this particular situation (for particular patient) where the likelihood threshold has been increased the type of behavior under consideration. Also "muscle fasciculat advances a lot because of its high likelihood compared to the contingencies. In the liability conscious behavior, the agent considers almost all consequences, except "message from PACU" because of its very long ti response which should allow for replanning (here "meteor strikes op room" becomes very high priority, since once it is considered - regardl its much too short response time allowed - its very high time pressur consequences make it very high priority. Similar arguments can be ma the results of each of the behavior models used in this demonstration.

The interpretation of our results shows (in the expert's opinion) they are reasonable and consistent with the generally acce] (execution-time) definition of each behavior model, and that there plausible explanation for the results that maps them into the corresp (conceptual) behaviors. These demonstrations again show that our for may at least provide a reasonable basis for representing and excha information and ideas about reaction-related behavior models, and th interpreting and studying different behaviors, in a considerable vari domains (from mundane tasks like car driving, to highly specialized on medical domains). A possible use is to start from a specific behavior (or the set of contingencies) exhibited by an agent, discover - using ma learning techniques - the parameters of the behavior model which er this behavior in our framework, and then use these parameter characterize the behavior and maybe to attempt to consciously mod However, these are only speculations at this point, since as stated before research is still needed to refine such a behavior description formalism useful tool for changing ideas among behavioral experts.

Appendix 4
Intensive Care Domain Experiments

We present here some of the results of the experiments we h conducted with our framework in the intensive care monitoring domai appendix mainly complements section 6.4.

#	Contingency (Response would be the typical response for this event)	Response time (min)	Conse-quences	Side-effects	Likeli-hood
1	myocardial-depression-post-cpb	10	8.5	7	3
2	myocardial-depression-sepsis	20	8	7.5	1
3	decreased-preload	20	7	3	7
4	increased-afterload	20	6.5	5	4
5					

2

cardiac-tamponade
5
8.5
7.5
3

6
hypovolemia
20
7
3
7

7
myocardial-ischemia
5
8
6
3

8
myocardial-infarction
60
6
5
3

9
right-heart-failure
10
8
7
2

10
digitalis-toxicity
180
5
4
2

11
postop-hypertension
20
6.5
5
4

12
cardiac-arrest
1

3

10
8
1

13
ventricular-fibrillation
1
10
8
1

14
ventricular-ectopy
5
7
7
6

15
sinus-bradycardia
5
7
5
3

16
atrial-fibrillation
20
7
6
4

17
paroxysmal-supraventric-tachycardia
20
6
6
4

18
ventricular-tachycardia
1
9
7
2

19
sinus-tachycardia
10
6
5

4

7

20
hypoxia
5
8
6
4

21
respiratory-acidosis
60
6
4
4

22
cardiogenic-pulmonary-edema
10
8.5
7
3

23
noncardiogenic-pulmonary-edema
20
8.5
8
2

Table A4.1 Contingencies for the ICU domain

24
atelectasis
120
6.5
5
6.5

25
pneumothorax
10
8
7
3

26
hemothorax
10
7
7

5

4

27
chylothorax
120
7
7
2

28
aspiration-pneumonia
240
8
5
1

29
pneumonia
240
7
5
3

30
diaphragmatic-paralysis
600
8
7
1

31
bronchospasm
30
8
7
4

32
pulmonary-embolism
10
8.5
7.5
3

33
ARDS
120
8.5
8
2

6

34
et-tube-disconnection
2
10
2
4

35
kinked-et-tube
5
8
2
4

36
right-mainstem-intubation
20
6.5
3
2

37
disseminated-intravascular-coagulat
60
8
7
2

38
dilutional-coagulopathy
60
7
3
5

39
platelet-deficiency
60
7
3
5

40
acute-hemolytic-transfusion-react
10
8.5
5
1

41
febrile-nonhemolytic-transfus-react

7

20
6.5
4
2

42
mechanical-bleeding
20
7.5
7.5
4

43
fibrinogen-defects
60
7
3
5

44
extrinsic-pathway-defects
60
7
3
5

45
intrinsic-pathway-defects
60
7
3
5

46
cerebrovascular-ischemia
60
8.5
7.5
2

47
cerebrovascular-embolism
30
9
7.5
1

48
endotoxemia
120
8.5

8

8
1

49
rewarming
240
3
3
7

50
hypothermia
240
4
4
7

51
hyperglycemia
120
5
4
2

52
metabolic-acidosis
60
6.5
4
3

53
acute-renal-failure
300
9
8
1

54
acute-tubular-necrosis
300
9
8
1

55
prerenal-azotemia
300
5
5
3

56
renal-azotemia
300
5
6
1

57
renal-embolism
300
7
7
1

58
high-cl
120
6
4
6

59
low-cl
120
6
4
2

60
high-ca
60
7
6
1

61
low-ca
60
6
3
6

62
low-mg
60
7
3
7

63

10

high-mg

60
8
5
2

64
low-na

30
7
2
2

65
high-na

60
6
3
2

66
dilutional-low-na

30
7
2
2

67
low-k

30
7.5
5
5

68
high-k

30
8
7
4

Table A4.1 Contingencies for the ICU domain (continued)

#
Contingency (Response would be the typical response for this event)
Resp.
time
Conse-
quences
Side-
eff.

11

Likeli-
hood
Criti-
cality

34
et-tube-disconnection
2
10
2
4
4.2E12

18
ventricular-tachycardia
1
9
7
2
2.2E12

13
ventricular-fibrillation
1
10
8
1
6.1E11

12
cardiac-arrest
1
10
8
1
6.1E11

35
kinked-et-tube
5
8
2
4
1.8E10

20
hypoxia
5
8
6
4

12

2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

1

13

myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism
10
8.5
7.5
3
2.13E7

6
hypovolemia
20
7
3
7
2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax
10
8
7
3
2.01E7

40
acute-hemolytic-transfusion-react
10
8.5
5
1
1.28E7

26
hemothorax
10
7

14

7
4
1.05E7

9
right-heart-failure
10
8
7
2
8.94E6

11
postop-hypertension
20
6.5
5
4
1.38E6

4
increased-afterload
20
6.5
5
4
1.38E6

36
right-mainstem-intubation
20
6.5
3
2
1.23E6

16
atrial-fibrillation
20
7
6
4
9.78E5

41
febrile-nonhemolytic-transfus-react
20
6.5
4
2
6.98E5

15

67
low-k
30
7.5
5
5
6.63E5

42
mechanical-bleeding
20
7.5
7.5
4
3.54E5

66
dilutional-low-na
30
7
2
2
3.48E5

64
low-na
30
7
2
2
3.48E5

17
paroxysmal-supraventric-tachycardia
20
6
6
4
2.83E5

23
noncardiogenic-pulmonary-edema
20
8.5
8
2
1.81E5

68
high-k

16

30
8
7
4
1.47E5

31
bronchospasm
30
8
7
4
1.47E5

62
low-mg
60
7
3
7
8.57E4

45
intrinsic-pathway-defects
60
7
3
5
4.37E4

44
extrinsic-pathway-defects
60
7
3
5
4.37E4

43
fibrinogen-defects
60
7
3
5
4.37E4

39
platelet-deficiency
60
7
3

17

5
4.37E4

38
dilutional-coagulopathy
60
7
3
5
4.37E4

2
myocardial-depression-sepsis
20
8
7.5
1
4.26E4

61
low-ca
60
6
3
6
3.21E4

47
cerebrovascular-embolism
30
9
7.5
1
1.58E4

21
respiratory-acidosis
60
6
4
4
7.63E3

Table A4.2. ICU domain contingencies ordered by criticality for T_{min} = 0.5 (2 hours) and L_{min} = 1

52
metabolic-acidosis
60
6.5
4

18

3
6.46E3

63
high-mg
60
8
5
2
4.76E3

65
high-na
60
6
3
2
3.57E3

8
myocardial-infarction
60
6
5
3
1.94E3

46
cerebrovascular-ischemia
60
8.5
7.5
2
1.22E3

37
disseminated-intravascular-coagulat
60
8
7
2
1.14E3

58
high-cl
120
6
4
6
5.36E2

19

24
atelectasis
120
6.5
5
6.5
4.70E2

60
high-ca
60
7
6
1
2.51E2

59
low-cl
120
6
4
2
59.63

33
ARDS
120
8.5
8
2
23.32

51
hyperglycemia
120
5
4
2
22.46

27
chylothorax
120
7
7
2
10.64

48
endotoxemia
120

20

8.5
8
1
5.83

29
pneumonia
240
7
5
3
2.21

10
digitalis-toxicity
180
5
4
2
1.71

50
hypothermia
240
4
4
7
1.52

49
rewarming
240
3
3
7
1.32

28
aspiration-pneumonia
240
8
5
1
1.07

55
prerenal-azotemia
300
5
5
3

21

0.41

54
acute-tubular-necrosis
300
9
8
1
0.32

53
acute-renal-failure
300
9
8
1
0.32

57
renal-embolism
300
7
7
1
0.16

56
renal-azotemia
300
5
6
1
5.9E-2

30
diaphragmatic-paralysis
600
8
7
1
5.3E-2

Table A4.2. ICU domain contingencies ordered by criticality for T_{min} = 0.5 (2 hours) and L_{min} = 1 (continued)

Table A4.1 lists the entire set of 68 contingencies defined by the e in the domain for the situations described in figure 6.1, together witl characteristic valuesm. The contingencies are listed in the order specif the experts (grouped by categories of complications that may develop).

The first part of this demonstration consisted in running the crit function part of the framework on this data set, for the recomm behavior model (section 6.3). We have done this for several expert n which differ in the minimum time pressure threshold value (T_{min}), and the minimum likelihood threshold (L_{min}) value. We shall present here only the results of four such experiments, although we have made a much l number.

Table A4.2 shows the order of the contingencies given by the "nor behavior model for a maximum reaction time of 2 hours (T_{min} = 0.50) and a minimum likelihood of 1. The rest of the expert model parameters ar unchanged during all these experiments (they are: f_{tc} = 60 / $time_{rc}$; T_{ma} (36 seconds); CS_{min} = 2.3).

#	Contingency (Response would be the typical response for this event)	Resp. time	Conse- quences	Side- eff.	Likeli- hood	Criti- cality
34	et-tube-disconnection	2	10	2	4	4.2E12
18	ventricular-tachycardia	1	9	7	2	2.2E12
35	kinked-et-tube	5	8	2	4	1.8E10
20	hypoxia	5				

23

8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3

24

3.26E7

1
myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism
10
8.5
7.5
3
2.13E7

6
hypovolemia
20
7
3
7
2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax
10
8
7
3
2.01E7

26
hemothorax
10
7
7
4
1.05E7

9

25

right-heart-failure
10
8
7
2
8.94E6

11
postop-hypertension
20
6.5
5
4
1.38E6

4
increased-afterload
20
6.5
5
4
1.38E6

36
right-mainstem-intubation
20
6.5
3
2
1.23E6

16
atrial-fibrillation
20
7
6
4
9.78E5

13
ventricular-fibrillation
1
10
8
1
7.86E5

12
cardiac-arrest
1
10

26

8
1
7.86E5

41
febrile-nonhemolytic-transfus-react
20
6.5
4
2
6.98E5

67
low-k
30
7.5
5
5
6.63E5

42
mechanical-bleeding
20
7.5
7.5
4
3.54E5

66
dilutional-low-na
30
7
2
2
3.48E5

64
low-na
30
7
2
2
3.48E5

17
paroxysmal-supraventric-tachycardia
20
6
6
4
2.83E5

23
noncardiogenic-pulmonary-edema
20
8.5
8
2
1.81E5

68
high-k
30
8
7
4
1.47E5

31
bronchospasm
30
8
7
4
1.47E5

62
low-mg
60
7
3
7
8.57E4

45
intrinsic-pathway-defects
60
7
3
5
4.37E4

44
extrinsic-pathway-defects
60
7
3
5
4.37E4

43
fibrinogen-defects

28

60
7
3
5
4.37E4

39
platelet-deficiency
60
7
3
5
4.37E4

38
dilutional-coagulopathy
60
7
3
5
4.37E4

61
low-ca
60
6
3
6
3.21E4

21
respiratory-acidosis
60
6
4
4
7.63E3

Table A4.3. ICU domain contingencies ordered by criticality for T_{min} = 0.5 (2 hours) and L_{min} = 2

52
metabolic-acidosis
60
6.5
4
3
6.46E3

63
high-mg

29

60
8
5
2
4.76E3

40
acute-hemolytic-transfusion-react
10
8.5
5
1
3.59E3

65
high-na
60
6
3
2
3.57E3

8
myocardial-infarction
60
6
5
3
1.94E3

46
cerebrovascular-ischemia
60
8.5
7.5
2
1.22E3

37
disseminated-intravascular-coagulat
60
8
7
2
1.14E3

58
high-cl
120
6
4

30

6
5.36E2

24
atelectasis
120
6.5
5
6.5
4.70E2

2
myocardial-depression-sepsis
20
8
7.5
1
2.06E2

47
cerebrovascular-embolism
30
9
7.5
1
1.25E2

59
low-cl
120
6
4
2
59.63

33
ARDS
120
8.5
8
2
23.32

51
hyperglycemia
120
5
4
2
22.46

31

60
high-ca
60
7
6
1
15.86

27
chylothorax
120
7
7
2
10.64

48
endotoxemia
120
8.5
8
1
2.41

29
pneumonia
240
7
5
3
2.21

10
digitalis-toxicity
180
5
4
2
1.71

50
hypothermia
240
4
4
7
1.52

49
rewarming
240

32

3
3
7
1.32

28
aspiration-pneumonia
240
8
5
1
1.07

55
prerenal-azotemia
300
5
5
3
0.41

54
acute-tubular-necrosis
300
9
8
1
0.32

53
acute-renal-failure
300
9
8
1
0.32

57
renal-embolism
300
7
7
1
0.16

56
renal-azotemia
300
5
6
1

33						5.9E-2
30	diaphragmatic-paralysis	600	8	7	1	5.3E-2

Table A4.3. ICU domain contingencies ordered by criticality for T_{min} = 0.5 (2 hours) and L_{min} = 2 (continued)

To show the effect of varying the likelihood parameter in the e model, table A4.3 presents the ordering of contingencies according to th behavior model, with all the parameters unchanged except the mir likelihood raised at 2. We can see that highly consequential but low like contingencies like ventricular-fibrillation and cardiac-arrest experience a significant drop in criticality (from the 3rd place to the 22nd). Howeve high consequences and high time pressure ensure that they do not fa much (they are still ranked by the framework in the first third of a contingencies considered).

#	Contingency (Response would be the typical response for this event)	Resp. time	Conse-quences	Side-eff.	Likeli-hood	Criti-cality
34	et-tube-disconnection	2	10	2	4	4.2E12
18	ventricular-tachycardia	1	9	7	2	2.2E12

34

13
ventricular-fibrillation
1
10
8
1
6.1E11

12
cardiac-arrest
1
10
8
1
6.1E11

35
kinked-et-tube
5
8
2
4
1.8E10

20
hypoxia
5
8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5

35

7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
8.21E7

22
cardiogenic-pulmonary-edema
10
8.5
7
3
3.26E7

1
myocardial-depression-post-cpb
10
8.5
7
3
3.26E7

32
pulmonary-embolism
10
8.5
7.5
3
2.13E7

6
hypovolemia
20
7
3
7

36

2.08E7

3
decreased-preload
20
7
3
7
2.08E7

25
pneumothorax
10
8
7
3
2.01E7

40
acute-hemolytic-transfusion-react
10
8.5
5
1
1.28E7

26
hemothorax
10
7
7
4
1.05E7

9
right-heart-failure
10
8
7
2
8.94E6

11
postop-hypertension
20
6.5
5
4
1.38E6

4

37

increased-afterload
20
6.5
5
4
1.38E6

36
right-mainstem-intubation
20
6.5
3
2
1.23E6

16
atrial-fibrillation
20
7
6
4
9.78E5

41
febrile-nonhemolytic-transfus-react
20
6.5
4
2
6.98E5

67
low-k
30
7.5
5
5
6.63E5

42
mechanical-bleeding
20
7.5
7.5
4
3.54E5

66
dilutional-low-na
30
7

38

2
2
3.48E5

64
low-na
30
7
2
2
3.48E5

17
paroxysmal-supraventric-tachycardia
20
6
6
4
2.83E5

23
noncardiogenic-pulmonary-edema
20
8.5
8
2
1.81E5

68
high-k
30
8
7
4
1.47E5

31
bronchospasm
30
8
7
4
1.47E5

2
myocardial-depression-sepsis
20
8
7.5
1
4.26E4

47
cerebrovascular-embolism
30
9
7.5
1
1.58E4

62
low-mg
60
7
3
7
2.92E2

45
intrinsic-pathway-defects
60
7
3
5
2.09E2

44
extrinsic-pathway-defects
60
7
3
5
2.09E2

43
fibrinogen-defects
60
7
3
5
2.09E2

39
platelet-deficiency
60
7
3
5
2.09E2

38
dilutional-coagulopathy

40		60	7	3	5	2.09E2
61	low-ca	60	6	3	6	1.79E2
21	respiratory-acidosis	60	6	4	4	87.36

Table A4.4. ICU domain contingencies ordered by criticality for T_{min} = 2 (30 minutes) and L_{min} = 1

52	metabolic-acidosis	60	6.5	4	3	80.43
63	high-mg	60	8	5	2	69.02
65	high-na	60	6	3	2	59.77
8	myocardial-infarction	60				

41

6
5
3
44.05

46
cerebrovascular-ischemia
60
8.5
7.5
2
34.95

37
disseminated-intravascular-coagulat
60
8
7
2
33.91

58
high-cl
120
6
4
6
23.16

24
atelectasis
120
6.5
5
6.5
21.70

60
high-ca
60
7
6
1
15.86

59
low-cl
120
6
4
2

42

7.72

33
ARDS

120
8.5
8
2
4.82

51
hyperglycemia

120
5
4
2
4.73

27
chylothorax

120
7
7
2
3.26

48
endotoxemia

120
8.5
8
1
2.41

29
pneumonia

240
7
5
3
2.21

10
digitalis-toxicity

180
5
4
2
1.71

50

43

hypothermia
240
4
4
7
1.52

49
rewarming
240
3
3
7
1.32

28
aspiration-pneumonia
240
8
5
1
1.07

55
prerenal-azotemia
300
5
5
3
0.41

54
acute-tubular-necrosis
300
9
8
1
0.32

53
acute-renal-failure
300
9
8
1
0.32

57
renal-embolism
300
7

7
1
0.16

56
renal-azotemia
300
5
6
1
5.9E-2

30
diaphragmatic-paralysis
600
8
7
1
5.3E-2

Table A4.4. ICU domain contingencies ordered by criticality for T_{min} = 2 (30 minutes) and L_{min} = 1 (continued)

Tables A4.4 and A4.5 show the effect of increasing the time pres threshold. While table A4.2 contains the contingencies ordered according expert model which recommends reactions for contingencies with al response time of up to 2 hours from the time a contingency is detectec A4.4 reduces this time to half an hour (minimum time pressure T_{min} = table A4.5 reduces it even further, to just 5 minutes (minimum time p T_{min} = 12). Notice that contingencies with very low likelihood but highe pressure (like myocardial-depression-sepsis and cerebrovascular-embo advance over more likely contingencies but with time pressure lower th recommended reaction threshold, in table A4.4. However, when the pressure threshold is raised significantly more (table A4.5), we obta identical ordering with the initial one in table A4.2, because the expe recommended reactions only for very time critical contingencies, which ranked as having high criticality by the framework even from the beg other things being equal. There is however a significant difference be tables A4.2 and A4.5 (and to a lesser extent table A4.4), namely a threshold for monitoring. In the case of a very low time pressure thresh hours), there is no such clear threshold, since the criticality of contin decreases gradually in table A4.2, without a clear gap. This is because, the maximum reaction time recommended is very large, the time pressu contingencies with long allowed response time is so small anyway, that i not influence the criticality of that contingency too much. This contrast the cases when the maximum reaction time recommended is small, for the time pressure is high enough to make a significant difference i criticality value. This is why in table A4.5 we have a clear threshold (giv a significant gap in the sequence of criticality values) after the contingency in the sequence (cardiac-tamponade). The same pheno takes place in table A4.4 after the cerebrovascular-embolism contingenc

#	Contingency (Response would be the typical response for this event)	Resp. time	Conse-quences	Side-eff.	Likeli-hood	Criti-cality
34	et-tube-disconnection	2	10	2	4	4.2E12
18	ventricular-tachycardia	1	9	7	2	2.2E12
13	ventricular-fibrillation	1	10	8	1	6.1E11
12	cardiac-arrest	1	10	8	1	6.1E11
35	kinked-et-tube	5	8	2	4	

46

1.8E10

20
hypoxia
5
8
6
4
2.53E9

7
myocardial-ischemia
5
8
6
3
1.42E9

15
sinus-bradycardia
5
7
5
3
1.24E9

14
ventricular-ectopy
5
7
7
6
7.62E8

5
cardiac-tamponade
5
8.5
7.5
3
6.84E8

19
sinus-tachycardia
10
6
5
7
9.06E3

22

47

cardiogenic-pulmonary-edema
10
8.5
7
3
5.71E3

1
myocardial-depression-post-cpb
10
8.5
7
3
5.71E3

32
pulmonary-embolism
10
8.5
7.5
3
4.62E3

6
hypovolemia
20
7
3
7
4.56E3

3
decreased-preload
20
7
3
7
4.56E3

25
pneumothorax
10
8
7
3
4.48E3

40
acute-hemolytic-transfusion-react
10
8.5

48

5
1
3.59E3

26
hemothorax
10
7
7
4
3.25E3

9
right-heart-failure
10
8
7
2
2.99E3

11
postop-hypertension
20
6.5
5
4
1.17E3

4
increased-afterload
20
6.5
5
4
1.17E3

36
right-mainstem-intubation
20
6.5
3
2
1.11E3

Table A4.5. ICU domain contingencies ordered by criticality for T_{min} = 12 (5 minutes) and L_{min} = 1

16
atrial-fibrillation
20
7

49

6
4
9.88E2

41
febrile-nonhemolytic-transfus-react
20
6.5
4
2
8.35E2

67
low-k
30
7.5
5
5
8.14E2

42
mechanical-bleeding
20
7.5
7.5
4
5.95E2

66
dilutional-low-na
30
7
2
2
5.90E2

64
low-na
30
7
2
2
5.90E2

17
paroxysmal-supraventric-tachycardia
20
6
6
4
5.32E2

50

23
noncardiogenic-pulmonary-edema
20
8.5
8
2
4.25E2

68
high-k
30
8
7
4
3.83E2

31
bronchospasm
30
8
7
4
3.83E2

62
low-mg
60
7
3
7
2.92E2

45
intrinsic-pathway-defects
60
7
3
5
2.09E2

44
extrinsic-pathway-defects
60
7
3
5
2.09E2

43
fibrinogen-defects

51

60
7
3
5
2.09E2

39
platelet-deficiency
60
7
3
5
2.09E2

38
dilutional-coagulopathy
60
7
3
5
2.09E2

2
myocardial-depression-sepsis
20
8
7.5
1
2.06E2

61
low-ca
60
6
3
6
1.79E2

47
cerebrovascular-embolism
30
9
7.5
1
1.25E2

21
respiratory-acidosis
60
6
4

52

4
87.36

52
metabolic-acidosis
60
6.5
4
3
80.43

63
high-mg
60
8
5
2
69.02

65
high-na
60
6
3
2
59.77

8
myocardial-infarction
60
6
5
3
44.05

46
cerebrovascular-ischemia
60
8.5
7.5
2
34.95

37
disseminated-intravascular-coagulat
60
8
7
2
33.91

53

58
high-cl
120
6
4
6
23.16

24
atelectasis
120
6.5
5
6.5
21.70

60
high-ca
60
7
6
1
15.86

59
low-cl
120
6
4
2
7.72

33
ARDS
120
8.5
8
2
4.82

51
hyperglycemia
120
5
4
2
4.73

27
chylothorax
120

54

7
7
2
3.26

48
endotoxemia

120
8.5
8
1
2.41

29
pneumonia

240
7
5
3
2.21

10
digitalis-toxicity

180
5
4
2
1.71

50
hypothermia

240
4
4
7
1.52

49
rewarming

240
3
3
7
1.32

28
aspiration-pneumonia

240
8
5
1

55

1.07

55
prerenal-azotemia
300
5
5
3
0.41

54
acute-tubular-necrosis
300
9
8
1
0.32

53
acute-renal-failure
300
9
8
1
0.32

57
renal-embolism
300
7
7
1
0.16

56
renal-azotemia
300
5
6
1
5.9E-2

30
diaphragmatic-paralysis
600
8
7
1
5.3E-2

Table A4.5. ICU domain contingencies ordered by criticality for T_{min} = 12 (5 minutes) and L_{min} = 1 (continued)

The most important conclusion to be drawn from this demonstrati that the recommendations of our framework were found to be reasona our domain experts. They have agreed, in each case (i.e. for each expert used) with the ordering of the contingencies proposed by our system, them reasonable and finding reasonable interpretations for them. Sinc is no other (objective) way to evaluate the framework's recommendatio may conclude that the framework and the "normal" behavior model w defined are a reasonable solution to our original problem.

References

[Agre & Chapman, 1987] Agre,P.E., Chapman,D., Pengi: An Implementation Theory of Activity, in Proceedings of AAAI-1987, pp. 268-272.

[Ash & Hayes-Roth, 1993] Ash,D., Hayes-Roth,B., A Comparison of Action-Ba Hierarchies and Decision Trees for Real-Time Performanc in Proceedings of AAAI-1993.

[Ash, 1994] Ash,D., Diagnosis Using Action-Based Hierarchies for Opti Real-time Performance, Ph.D. Dissertation, Computer Scien Department, Stanford Univ., Stanford, CA, 1994.

[Brooks, 1982] Brooks,R.A., Symbolic Error Analysis and Robot Programm in Int. Journal of Robotics Research, Vol. 4, 1982, pp.29-68

[Brooks, 1986] Brooks, R.A., A Robust Layered Control System for a Mo Robot, in IEEE Journal of Robotics and Automation, RA-2, pp.14-23, April, 1986.

[Brooks, 1991] Brooks, R.A., Intelligence Without Representation, in Artif Intelligence Vol.47, 1991, pp.139-160.

[Chapman, 1987] Chapman,D., Planning for Conjunctive Goals, in Artif Intelligence Vol.32, 1987, pp.333-337.

[Cohen & al., 1989] Cohen,P.R., Greenberg,M., Hart,D.M., Howe,A.E., Trial By Understanding the Design Requirements for Agents in Comp Environments, in AI Magazine 10:3, 1989, pp.32-48.

[Dabija, 1990] Dabija,V.G., Learning by Experimentation in Uncer Environments, in Proceedings of the 4-th Australian Joint Confer on Artificial Intelligence, AI'90, Perth, Australia, pp.24-37, 1990.

[Dabija & al., 1992a] Dabija,V.G., Tsujino,K., Nishida,S., Learning to Le Decision Trees, in Proceedings of the 10-th National Conference Artificial Intelligence, AAAI-92, San Jose, 1992, pg.88-95.

[Dabija & al., 1992b] Dabija,V.G., Tsujino,K., Nishida,S., "Theory Formation i Decision Trees Domain", in Journal of the Japanese Society for Artificial Intelligence, vol.7, no.3, May 1992, pg.518-529.

[DeAnda & Gaba, 1991] DeAnda,A., Gaba,D.M., Role of Experience in Response to Simulated Critical Incidents, in Anesthesia and Analg No. 72, 1991, pg. 308-315.

[Drummond, 1989] Drummond,M.E., Situated Control Rules, in Proceedings of the First International Conference on Principles of Knowled Representation and Reasoning, Toronto, 1989, pp.103-113.

[Dufay and Latombe, 1984] Dufay,B., Latombe,J.C., An Approach to Auto Robot Programming Based on Inductive Learning, Int. Journal of Robotics Research, Vol. 3, No.4, 1984, pp.3-20.

[FAA, 1991] Federal Aviation Administration, Aeronautical Decision Making, Advisory Circular AC-60-22, December 13, 1991.

[Fikes & Nilsson, 1971] Fikes,R.E., Nilsson,N.J., STRIPS : A New Approach to Application of Theorem Proving to Problem Solving, in Artific Intelligence Vol.2, No.3, 1971, pp.189-208.

[Fikes & al., 1972] Fikes,R.E., Hart,P.E., Nilsson,N.J., Learning and Execut Generalized Robot Plans, in Artificial Intelligence Vol.3, No.4, 1972, pp.251-288.

[Firby, 1987] Firby,R.J., An Investigation into Reactive Planning in Com Domains, in Proceedings of AAAI-1987, pp.202-206.

[Fish & al., 1991] Fish,K.J., Gaba,D.M., Howard,S.K., Sarnquist,F.H., Catalog Critical Incidents in Anesthesiology, 1991.

[Gaba & al., 1991] Gaba,D.M., Fish,K.J., Howard,S.K., Principles of Anesth Crisis Resource Management, in Anesthesia Crisis Resourc Management, Training Program Course Material, Stanford Univ., 19

[Gaba, 1991] Gaba,D.M., Pre Case Planning, in Pre-anesthesia Crisis Resc Management, Training Program Course Material, Stanford Univ., 19

[Georgeff & Lansky, 1987] Georgeff,M.P., Lansky, A.L., Reactive Reasoning Planning, in Proceedings of AAAI-1987, pp.677-682.

[Georgeff, 1989] Georgeff,M.P., An Embedded Reasoning and Planning Sy in Proceedings of the Rochester Planning Workshop: From Form Systems to Practical Systems, Univ. of Rochester, 1989, pp.105-128.

[Ginsberg, 1989] Ginsberg,M.L., Universal Planning: An (Almost) Univers Bad Idea, in A.I. Magazine, Vol.10, Nr.4, Winter 1989, pp.40-44.

[Hayes-Roth, 1985] Hayes-Roth,B., A Blackboard Architecture for Contro Artificial Intelligence, Vol.26, pp.251-321, 1985.

[Hayes-Roth, 1990] Hayes-Roth,B., Architectural Foundations of Real-T Performance in Intelligent Agents, in The Journal of Real-Ti Systems, Vol.2, pg.99-125, 1990.

[Hayes-Roth & al., 1992] Hayes-Roth,B., Washington,R., Ash,D., Hewett, Collinot,A., Vina,A., Seiver,A., Guardian: a Prototype Intelligent Ag for Intensive-Care Monitoring, Artificial Intelligence in Medicine Vol.4, pg.165-185, 1992.

[Hayes-Roth, 1993] Hayes-Roth, B., Opportunistic Control of Action Intelligent Agents, in IEEE Transactions on Systems, Man, ar Cybernetics, Special Issue on Planning, Scheduling, and Control: AI-based Integrated Systems View, March, 1993.

[Hendler & Agrawala, 1990] Hendler,J., Agrawala,A., Mission Critical Plan AI on the MARUTI Real-Time Operating System, in Proceedings of t DARPA Workshop on Innovative Approaches to Planning, Schedul and Control, 1990, pp.77-84.

[Horvitz, 1989] Horvitz,E.J., Rational Metareasoning and Compilation Optimizing Decisions under Bounded Resources, in Proceedings of Computational Intelligence, ACM, Milan, Italy, September 1989.

[Hutchinson and Kak, 1990] Hutchinson,S.A., Kak,A.C., SPAR: A Planner Satisfies Operational and Geometric Goals in Uncertain Environment AI Magazine, Vol.11, No.1, Spring 1990, pp.30-61

[Kaelbling, 1987] Kaelbing,L.P., An Architecture for Intelligent Reac Systems, in M.Georgeff and A.Lansky, (Eds), Reasoning About Action and Plans: Proceedings of the 1986 Workshop, Morgan Kaufmann,19

[Laird & Rosenbloom, 1990] Laird,J.E., Rosenbloom,P.S., Integrating Execu Planning, and Learning in Soar for External Environment inProceedings of AAAI-1990, pp. 1022-1029.

[Latombe & al., 1991] Latombe,J.C., Lazanas,A., Shekhar,S., Robot Mot
Planning with Uncertainty in Control and Sensing, Artificial
Intelligence, Vol. 52 (1991) pp.1-47
[Lozano-Perez, 1976] Lozano-Perez,T., The Design of a Mechanical Asser
System, Tech.Rep.AI-TR 397, AILab., MIT, Cambridge, MA, 1976
[Marsh & Greenwood, 1986] Marsh,J., Greenwood,J., Real-time AI: softv
architecture issues, in Proceedings of the IEEE National Aerospace and
Electronics Conference, pp.67-77, Washington, D.C., 1986.
[Martin & Allen, 1990] Martin,N.J., Allen,J.A., Combining Reactive and Stra
Planning through Decomposition Abstraction, in Proceedings of
DARPA Workshop on Innovative Approaches to Planning, Schedul
and Control, 1990, pp. 301-312.
[McCarthy, 1958] McCarthy, J., Programs with Common Sense, in M.Mir
(ed.), Semantic Information Processing, The MIT Press, Cambridge,
1968, pp.403-418; section1: The Advice Taker, is a reprint of a p
taken from Mechanization of Thought Processes, Vol.1, pp.77-84
Proceedings Symposium, National Physical Laboratory, Londo
November 24-27, 1958.
[McCarthy, 1977] McCarthy, J., Epistemological Problems of Artific
Intelligence, in Proceedings of the Fifth International Joint Confer
on Artificial Intelligence - IJCAI-77, pp.1038-1044, 1977.
[Minton, 1990] Minton,S., Quantitative Results Concerning the Utility
Explanation-Based Learning, Artificial Intelligence, Vol.42:1,
pp.363-392, 1990.
[Mitchell, 1978] Mitchell,T.M., Version Spaces: An Approach to Conc
Learning, PhD. Dissertation, Comp.Sci.Dept, Stanford U., Stanford,
1978.
[Mitchell & al., 1983] Mitchell, T.M., Utgoff,P.E., Banerji,R.B., Learning
Experimentation: Acquiring and Refining Problem Solving Heurist
in R.S.Michalski, J.G.Carbonell, T.M.Mitchell (eds.), Machine Learnir
An Artificial Intelligence Approach, Tioga, Palo Alto, CA, 1983,
161-190.
[Mitchell & al., 1986] Mitchell, T.M., Keller,R.M., Kedar-Cabelli,S.T
Explanation-Based Generalization: A Unifying View, Machine Learn
1:1, pp.47-80, 1986.
[Mitchell, 1990] Mitchell,T.M., Becoming Increasingly Reactive, in Proceedings
of AAAI-1990, pp.1051-1058.
[Musliner et. al., 1994] Musliner,D.J., Durfee,E.H., Shin,K.J., World Modeling
the Dynamic Construction of Real-Time Control Plans, Artifici
Intelligence, (Special Issue on Computational Theories of Interac
and Agency), 1994, to appear.
[Nilsson, 1984] Nilsson, N.J., editor, Shakey the Robot. Technical Note
Stanford Research Institute, 1984.
[Nilsson, 1988] Nilsson, N.J., Action Networks, in Proceedings of the Roch
Planning Workshop, pp.21-52, Univ. of Rochester, 1988.
[Nilsson, 1992] Nilsson, N.J., Toward Agent Programs with Circuit Semar
Technical Report No. STAN-CS-92-1412, Computer Science Dept., Stanf
University, January, 1992.
[Pertin-Troccaz and Puget, 1987] Pertin-Troccaz,J., Puget,P., Dealing
Uncertainty in Robot Planning Using Program Proving Techniques,

R.C.Bolles and B.Roth, eds., Robotics Research 4, MIT Press, Cambridge, MA, 1987, pp.455-466
[Sacerdoti, 1975] Sacerdoti,E.D., The Non-Linear Nature of Plans, in Proceedings of the Fourth International Joint Conference on Artificial Intelliger IJCAI-75, pp.206-214, 1975.
[Schoppers, 1987] Schoppers,M.J., Universal Plans for Reactive Robots Unpredictable Environments, in Proceedings of the Tenth Internat Joint Conference on Artificial Intelligence - IJCAI-87, pp.1039-1C 1987.
[Stefik, 1981] Stefik,M.J., Planning with Constraints, in Artificial Intelligence Vol. 16, 1981. pp.111-140.
[Taylor, 1976] Taylor,R.H., Synthesis of Manipulator Control Programs Task-Level Specifications, Ph.D. Dissertation, Computer Scien Department, Stanford Univ., Stanford, CA, 1976
[Utgoff, 1988] Utgoff,P.E., Shift of Bias for Inductive Concept Learning, R.S.Michalski, J.G.Carbonell, T.M.Mitchell (eds.), Machine Learning: An Artificial Intelligence Approach, Vol.2, Morgan Kaufmann, Los Altos, CA, 1986, pp.107-148.
[Woods & al., 1987] Woods, D.D., Roth, E.M., Pople, H,Jr., "Cognitive Environ Simulation: an AI System for Human Performance Assessmen U.S.Nuclear Regulatory Commission Report, NUREG-CR-4862 (2), 1987.
[Yamada, 1992] Yamada,S., Reactive Planning with Uncertainty of a Plai Proceedings of the Third Annual Conference on AI, Simulation Planning in High Autonomy Systems, Perth, Australia, 1992, pp.201-

1 In order to simplify the analysis for clarity of exposition, we have delibe excluded the conventional driver's wisdom case that a ball popping up in the s usually followed by a running child.

2 Again we stress that, in this work we study conscious forms of reaction, prep planning time and consciously taken, as opposed to precognitive types of reacti locomotion type reaction).

3 Hopefully not at the same time...

4 In a specific setting (domain, expert, reactive planner and executing agent) parameters can be automatically or interactively learned using paradigms like t proposed in [Dabija, 1990; Dabija & al., 1992a,b].

5 We have specifically excluded the conventional driver's wisdom case that a ball up in the street is usually followed by a running child.