ELSEVIER

Reliability Engineering and System Safety 75 (2002) 207–214

RELIABILITY ENGINEERING & SYSTEM SAFETY

www.elsevier.com/locate/ress

Integrating objective and subjective hazard risk in decision-aiding system design

J.K. Kuchar*, D.S. Walton, D.M. Matsumoto

Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Cambridge, MA 02139, USA

Abstract

A generalized model is presented to incorporate objective (hard) and subjective (soft) hazard information in automated decision-aiding systems. The model may be used with more than one hazard, of more than one type, in a given problem. Uncertainties in state measurements, dynamics, hazard extent, and hazard severity are included, as is the consideration of the fact that different operators may have different concepts of what is an acceptable or unacceptable risk. By examining the tradeoffs created by these uncertainties, appropriate decision thresholds can be selected. Using an aviation case study, information gained from observation of aircraft behavior in the presence of weather was used to develop a model of weather as a soft hazard. This information can then be used in a decision aid to provide feedback on route acceptability. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Alerting; Probabilistic modeling; Hazards; Decision-aids; Decision theory; Risk

1. Introduction

Real-time decision aiding and alerting systems are often used to assist human operators in controlling processes efficiently and in preventing undesirable incidents from occurring (such as a collision in a vehicle control application, or exceeding temperature limits in process control). There are many types of real-time decision-aids, ranging from process status displays, to planning tools, to safety- and time-critical warning systems. To date, warning systems have generally been restricted to cases in which there is a clear definition of hazardous states. For example, traffic collision risk can be defined in concrete, objective terms (e.g. no closer than 100 m separation between aircraft), which then is translated into algorithms and decision thresholds. This can be classified as a case of *objective* assessment of hazard risk. Due to sensor and prediction errors, there may still be uncertainty in whether a decision to change the process' trajectory is needed. These uncertainties, however, can also be objectively estimated and used when defining decision thresholds to balance false alarms and missed detections, and optimize system performance from the human operator's perspective.

In cases in which the distinction between hazard and non-hazard is less distinct (i.e. the hazard risk is *subjective*), decision-aids typically display the state information but leave the decision-making to the human operator. Aviation examples of subjective hazards include weather precipitation levels, turbulence intensities, forecast icing, or visibility.

For an automation tool to be accepted, the decisions and feedback it provides should be aligned with operator mental models and expectations. A decision aid that does not consider subjective hazards may generate inappropriate decisions that decrease operator confidence and acceptance. For example, several automation tools are being developed for detecting and resolving air traffic conflicts and managing the arrival flow at airports [1–3]. These tools currently do not include hazardous weather information in their automated decisions, though some study of incorporating weather has begun in this area [4]. When weather is not considered by the automation, the human operator must mentally integrate the information to determine whether a given automated suggestion is appropriate. This may increase workload and decrease the utility of the automation in poor weather conditions. There is an opportunity, then, to enhance the automation by including subjective information in its decision-making process. This paper describes a general modeling approach that integrates subjective and objective hazard risk into a form that can be used in an automated decision aid. This facilitates incorporating information about multiple hazard sources, both subjective (e.g. weather) and objective (e.g. traffic and terrain), into a decision aid or alerting system. A specific example application is also presented in which precipitation intensity information was used to develop a model of weather that

* Corresponding author. Tel.: +1-617-252-1512; fax: +1-617-253-4196.
E-mail address: jkkuchar@mit.edu (J.K. Kuchar).

PII: S0951-8320(01)00095-3

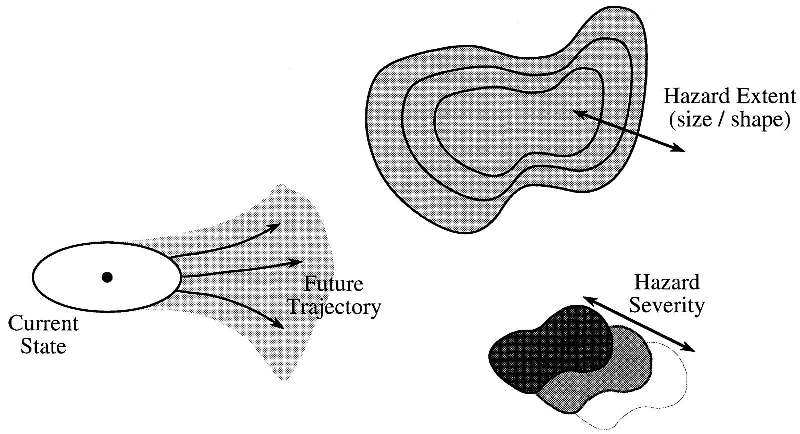

Fig. 1. Uncertainties in the decision to alert.

can be integrated with other hazard information to provide feedback on route acceptability.

2. General model

The class of decision-aiding problems considered in this paper involves those in which a process of interest is controlled through a combination of humans and automation in an environment in which undesirable hazardous states may exist. The operator's task is to control the process to arrive at some desired end state without experiencing an undesirable incident. Example processes may include chemical or power processes, single vehicles, large-scale transportation systems, financial markets, or other applications. The operator may be a single human or a combination of automation and humans.

If appropriate criteria are met, the decision-aiding (or alerting) system provides additional information to the operator with the intent of bringing the hazard to the operator's attention and, in some cases, to explicitly aid the operator in resolving the problem. In this way, the alerting system acts as an independent safety-enhancing system, which operates according to some predefined decision threshold logic.

2.1. Hazard encounters and incidents

To analyze alerting more completely, it is necessary to define the situations that must be avoided, termed here *incidents* and denoted by the event I. Example incidents include exceeding temperature or pressure limits that ruin a chemical process, the collision of two vehicles, or flight through severe turbulence. To facilitate the definition of incidents, it is first necessary to consider the process and environment in a state-space model. The appropriate choice of states depends on the application: in a vehicle control case, states could include position, velocity, and acceleration; in process control, states could include temperature, pressure, valve positions, and flow rates.

Because the occurrence of an incident may be a probabilistic event (e.g. flight through a region of heavy precipitation may involve severe turbulence, but it may not), the region in state-space in which an incident is possible is partitioned and is termed Hazard Space. Thus, entry into Hazard Space (termed a *hazard encounter*, or event E) is necessary but not sufficient for an incident to occur. The problem then reduces to one of determining whether Hazard Space will be encountered during the operation of the process, and if so, the likelihood for an incident will then result. Trajectories that do not penetrate Hazard Space can then be optimized to meet other constraints such as time, fuel burn, or other metrics of efficiency.

2.2. Uncertainties

With perfect information, entry into Hazard Space can be predicted exactly. Generally, however, this is not possible, due to the combination of four types of uncertainties. These uncertainties (current state, trajectory, extent, and severity) are shown schematically in Fig. 1 and discussed in more detail below.

First, the current state of the process may not actually be located at its estimated position in state-space. Current-state uncertainty (e.g. vehicle position or velocity) is typically a function of sensor errors. In some cases, the states of interest cannot be measured directly and must be inferred from other measurements; this may further increase uncertainty.

The second source of uncertainty relates to the projected future trajectory of the process. It is this projected trajectory, relative to Hazard Space, that is used to determine whether action is required at the current time to avoid an incident. In order to project the current state into the future, it is necessary to have a dynamic model of the process, the environment, and Hazard Space. Uncertainties in this model may exist regarding the dynamics of the controlled plant, sensors and actuators, the operator's actions, the environment, and the hazard itself. For example, aircraft motion can be accurately modeled using physics, but uncertainties in pilot behavior, winds, and navigational instrument errors may combine to produce an increasingly uncertain estimate of state position in time.

Uncertainty in the size, shape, or extent of Hazard Space leads to the third type of error in incident prediction. This type of uncertainty is specific to the hazard under consideration: collision hazards are generally well-defined (e.g. separation less than 152 m (500 ft) is considered to be a collision or 'near miss' in many aviation applications), while the boundaries of other types of hazards may be less certain (e.g. severe weather).

As discussed in Section 2.1, entry into Hazard Space may or may not result in an incident. The fourth type of uncertainty relates to the probability of an incident given where a hazard encounter has occurred. This can be thought of as a combination of an objective uncertainty in the severity of the hazard, and a subjective uncertainty in the definition of an incident. In the former, objective case, the structure of the hazard may be such that an incident occurs probabilistically following a hazard encounter: the hazard may be modeled using varying levels of 'hardness' or 'softness'. One example is flight over a missile site, which might have been destroyed previously; whether or not a missile is launched, can then be considered as a probabilistic event. While the estimation of this probability requires judgment, the threat posed by a missile launch could be considered objectively. In the latter, subjective uncertainty case, there may be differing opinion on what the proper definition of an incident is. Flight through poor weather, for example, may be acceptable to some operators and not to others. This acceptability is likely a function of many other factors such as operator experience and training, risk aversion or acceptance, the existence of alternate options, and the expected amount of time that will be spent in Hazard Space.

2.3. Decision tradeoffs

With hazards involving some uncertainty, any discrete decision to alert the operator or to otherwise determine whether a given trajectory is acceptable may be an error in one of the two ways. First, it may be the case that an alert was not necessary. Alternately, it may be the case that the decision to alert is never made or is made too late to prevent an incident (a missed or late alert). The tradeoff between these outcomes is a critical factor in designing an acceptable decision-aiding system and has been examined previously for objective hazards [5].

When dealing with cases in which the definition of an incident is subjective, an additional form of decision tradeoff occurs. Consider a case in which it is known with certainty that the process will enter a soft hazard. A decision to alert the operator may be an error if the operator does not consider the hazard sufficiently threatening, resulting in the perception of a false alarm. An analogue to a missed detection may instead occur if an alert is not made, but the operator would have desired an alert. This is not to say that an alerting system must always match an operator's mental model of when alerts should be issued — in some cases, the operator may have an incomplete concept of whether action is truly required. However, studies have shown the importance of designing systems to provide feedback so that operators can understand the reasoning or logic behind an automated decision [6].

While the decision tradeoff with an objective definition of an incident can be examined using models of dynamics and uncertainties, subjective incidents may require additional consideration of human factors, expert opinion, and operational experience. In some cases, it may be desirable to have operator-selectable or situation-dependent thresholds that can be tuned to the particular problem at hand. As one example, the threat posed by weather varies significantly depending on aircraft type, pilot experience, and other environmental factors such as overall extent of the storm, proximity to terrain, or the availability of escape routes. A study by MIT Lincoln Laboratory, for example, found that pilots were significantly more likely to penetrate severe weather as they went closer to the runway, possibly due to pressures of the constrained environment and to knowledge of the relatively short distance remaining to be flown [7].

The four types of uncertainty outlined above combine together with the result that whether an incident will occur can only be estimated with some probability. Section 3 develops a formal method for computing this probability. The remainder of the paper then discusses issues in hazard modeling for an aviation weather case study.

3. Computation of probability of incident

Fig. 2 depicts an example of the process state estimate, $\hat{\mathbf{x}}$, that includes some uncertainty. An error ellipse is shown in Fig. 2 that describes the region in state-space in which the true state $\mathbf{x}$ actually lies with some probability. Hazard Space, denoting a region where an incident can occur, is also shown.

Based on $\hat{\mathbf{x}}$, the alerting system must determine whether an alert is warranted. The need for an alert, however,

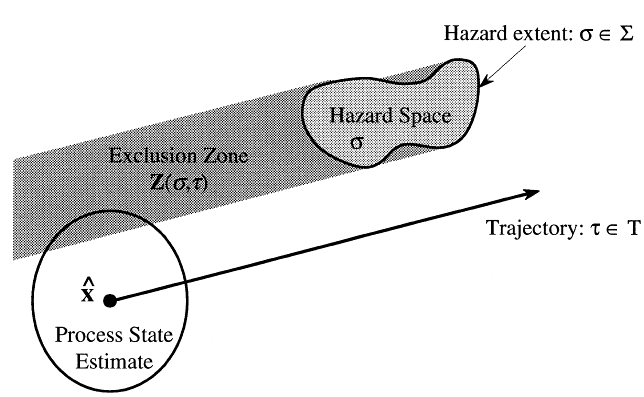

Fig. 2. State-space diagram.

depends on the trajectory that will be followed in the future. In Fig. 2, one possible trajectory is denoted τ. The actual trajectory that will be followed, however, may be any within some set T of trajectories. T depends on $\hat{\mathbf{x}}$, future control inputs, and knowledge of the system dynamics. Having defined T, and given a particular state estimate, $\hat{\mathbf{x}}$, there exists some probability that an incident will occur in the future, denoted $P(\mathrm{I}|\hat{\mathbf{x}})$. As discussed earlier, $P(\mathrm{I}|\hat{\mathbf{x}})$ is a function of uncertainties in the current state, future trajectory, hazard extent, and hazard severity.

Whether the decision-aiding system needs to alert the operator at the time shown in Fig. 2 depends on the value of $P(\mathrm{I}|\hat{\mathbf{x}})$. In general, the larger the probability that an incident will occur, the greater the need for an alert. An alert that is issued when $P(\mathrm{I}|\hat{\mathbf{x}})$ is small may be considered a false alarm if the human operator is aware of the hazard or would have avoided the hazard without the alert. However, if the alert is delayed until $P(\mathrm{I}|\hat{\mathbf{x}})$ is large, there may not be enough time or space to perform an avoidance maneuver and an incident may occur even if an alert is issued. A methodology to observe and design around this tradeoff has been previously developed [5].

As discussed earlier, an encounter with Hazard Space (event E) is necessary, but not sufficient for an incident (event I) to occur. Accordingly, $P(\mathrm{E}|\hat{\mathbf{x}})$ is defined as the probability that a hazard encounter occurs, and $P(\mathrm{I}|\hat{\mathbf{x}})$ is defined as the probability that an incident occurs in that same situation:

$$P(\mathrm{I}|\hat{\mathbf{x}}) = P(\mathrm{I}|\mathrm{E})P(\mathrm{E}|\hat{\mathbf{x}}) \tag{1}$$

Hazards for which $P(\mathrm{I}|\mathrm{E}) = 1$ are *hard hazards*. An encounter with a hard hazard means that an incident also occurs. An example hard hazard is a collision with another vehicle. Hazards for which $P(\mathrm{I}|\mathrm{E}) < 1$ are *soft hazards* — an encounter does not necessarily indicate that an incident will also occur. Soft hazards include severe weather or safety thresholds beyond which component failures may occur. Methods for modeling hard and soft hazards are discussed in Section 3.1.

Now, assume that the probability that the state is truly at some value $\mathbf{x}$ is given by the probability density function (PDF) $f_{\mathbf{x}}(\mathbf{x} - \hat{\mathbf{x}})$, typically based on sensor error distributions. Also, the size and shape of Hazard Space (defined by σ) is uncertain and may take on any value from a set of possible hazard extents, Σ, described by the PDF $f_\sigma(\sigma)$. The trajectory τ can take on any value from a set of possible trajectories, T, described by the PDF $f_\tau(\tau)$. Finally, an exclusion zone, $\mathbf{Z}(\sigma, \tau)$, can be defined to represent those locations where $\mathbf{x}$ must be for a hazard encounter to occur. Integrating $f_{\mathbf{x}}(\mathbf{x} - \hat{\mathbf{x}})$ over $\mathbf{Z}$ yields the probability that an encounter will occur for a given, known trajectory and hazard extent. Then, integrating over all possible trajectories and hazard extents gives the general probability of encounter:

$$P(\mathrm{E}|\hat{\mathbf{x}}) = \int_T \int_\Sigma \int_{\mathbf{Z}(\sigma,\tau)} f_{\mathbf{x}}(\mathbf{x} - \hat{\mathbf{x}}) f_\sigma(\sigma) f_\tau(\tau) \mathrm{d}x \, \mathrm{d}\sigma \, \mathrm{d}\tau \tag{2}$$

Provided that the PDFs are known or can be estimated, this expression can be solved using analytical methods, numerical integration, or Monte Carlo simulation [3,8,9]. Finally, $P(\mathrm{I}|\mathrm{E})$ is used in Eq. (1) to determine the probability of an incident, $P(\mathrm{I}|\hat{\mathbf{x}})$.

3.1. Hazard modeling

In cases where an encounter with Hazard Space is the same as an incident (hard hazards), $P(\mathrm{I}|\mathrm{E}) = 1$ and Eq. (2) directly yields the probability of such an incident occurring. When an encounter with a hazard does not necessarily mean that an incident will occur (a soft hazard), the methodology must be further modified to account for $P(\mathrm{I}|\mathrm{E}) < 1$. The probability of incident for soft hazards may be modeled in two ways. In the first, $P(\mathrm{I}|\mathrm{E})$ is independent of exposure. An example is a missile site that has a certain probability of being active on a particular day, regardless of how long the aircraft flies near it. The second model is one in which the probability of an incident depends on the time or distance over which an encounter occurs. An example is hazardous weather — the longer an aircraft remains inside a thunderstorm, the greater is the probability that an incident will occur.

When a soft time-independent hazard is encountered, $P(\mathrm{I}|\mathrm{E})$ is less than 1. Assuming that this probability is constant over Hazard Space, $P(\mathrm{I}|\hat{\mathbf{x}})$ can be determined using Eq. (1). When dealing with time-dependent hazards, $P(\mathrm{I}|\mathrm{E})$ depends on the amount of time or distance that the trajectory remains in Hazard Space. The total exposure to the hazard along each possible trajectory is then used to calculate $P(\mathrm{I}|\mathrm{E})$. This exposure can be estimated by integrating a hazard severity density function, $f_L(l)$, over the length of the trajectory through the hazard. For a constant-density threat, $f_L(l)$ can be represented by an exponential distribution. The integral of this PDF over a path of length L in the hazard results in:

$$P(\mathrm{I}|\mathrm{E}) = 1 - \mathrm{e}^{-L/\theta} \tag{3}$$

where θ represents the mean amount of exposure (in units of distance or time) until an incident occurs. As θ tends toward zero, the hazard becomes more like a hard hazard — very small exposures result in incidents. A large value of θ indicates an insubstantial hazard to which a large exposure is required before an incident will likely occur.

Because L is dependent on the trajectory and the hazard extent, the resulting equation for $P(\mathrm{I}|\hat{\mathbf{x}})$ for soft, time-dependent hazards is

$$P(\mathrm{I}|\hat{\mathbf{x}}) = \int_T \int_\Sigma \int_{\mathbf{Z}(\sigma,\tau)} (1 - \mathrm{e}^{-L(\sigma,\tau)/\theta}) f_{\mathbf{x}}(\mathbf{x} - \hat{\mathbf{x}}) f_\sigma(\sigma) f_\tau(\tau) \mathrm{d}x \, \mathrm{d}\sigma \, \mathrm{d}\tau \tag{4}$$

3.2. Multiple hazards

Additional complexity arises when more than one hazard may be encountered along a particular trajectory. Multiple hazards must be considered if the system is to integrate several potential threats when making decisions. Consider a situation in which there are two regions of Hazard Space, denoted A and B, along a single projected trajectory τ. Treating each hazard encounter as a separate event, the resultant probability of incident along the trajectory can be determined using:

$$P(\mathrm{I}_{\mathrm{A\ or\ B}}|\hat{\mathbf{x}}, \tau) = P(\mathrm{I}_{\mathrm{A}}|\hat{\mathbf{x}}, \tau) + P(\mathrm{I}_{\mathrm{B}}|\hat{\mathbf{x}}, \tau) - P(\mathrm{I}_{\mathrm{AB}}|\hat{\mathbf{x}}, \tau) \tag{5}$$

where the subscripts A and B indicate which hazard (or combination of hazards) is producing an incident, and τ is explicitly shown as a parameter to highlight the dependence on a specific trajectory. If the two hazards are conditionally independent of one another given the trajectory τ, this reduces to:

$$P(\mathrm{I}_{\mathrm{A\ or\ B}}|\hat{\mathbf{x}}, \tau) = P(\mathrm{I}_{\mathrm{A}}|\hat{\mathbf{x}}, \tau) + P(\mathrm{I}_{\mathrm{B}}|\hat{\mathbf{x}}, \tau) - P(\mathrm{I}_{\mathrm{A}}|\hat{\mathbf{x}}, \tau)P(\mathrm{I}_{\mathrm{B}}|\hat{\mathbf{x}}, \tau) \tag{6}$$

When more than one trajectory could be followed, Eq. (6) must be used on each trajectory τ separately, then integrated over the set of trajectories:

$$P(\mathrm{I}_{\mathrm{A\ or\ B}}|\hat{\mathbf{x}}) = \int_T [P(\mathrm{I}_{\mathrm{A}}|\hat{\mathbf{x}}, \tau) + P(\mathrm{I}_{\mathrm{B}}|\hat{\mathbf{x}}, \tau) - P(\mathrm{I}_{\mathrm{A}}|\hat{\mathbf{x}}, \tau)P(\mathrm{I}_{\mathrm{B}}|\hat{\mathbf{x}}, \tau)] f_\tau(\tau)\mathrm{d}\tau \tag{7}$$

In this way, a number of hazards, both hard and soft, can be assessed simultaneously.

4. Example application to aviation weather hazards

Up to this point, the focus has been on describing a method by which the probability of an incident can be computed, assuming that a model of the process and the hazards can be developed. This section discusses an aviation weather problem as a case study used to develop a model of a soft hazard.

The flight of an aircraft can be modeled generically as a process in which the pilot provides control inputs so as to arrive at some destination state. Currently, alerting systems are in place that warns pilots of hard hazard collision threats such as traffic or terrain. Pilots also have weather radar displays that depict precipitation intensity. Due to the soft, complex nature of weather as a hazard, pilots have traditionally had to integrate weather information with other constraints when determining tactical routes. As more complex alerting systems are developed, it may be attractive to incorporate soft weather information in the decision-aids, even if only at a fairly rudimentary level.

4.1. Weather and aircraft interaction data collection

Weather is a complex hazard, and translating weather information into a form that can be used by an automated system is a challenge that will continue to be addressed by researchers in the future. As a preliminary step in this direction, however, observations of enroute aircraft proximity to weather were performed to develop a simplified, prototype model of weather as a soft hazard.

Courtesy of the MIT Lincoln Laboratory, archived aircraft track and weather data were obtained for the hours between 2100 GMT on May 19, 1997 and 0900 GMT on May 20, 1997 from the Dallas Fort-Worth enroute sector, which spans approximately 1000 km from New Mexico across Texas [7]. The aircraft position information for the enroute airspace above 30,000 ft was updated every 6 s, and the weather precipitation data was updated every 5 min. A total of 1095 aircraft were included in the track data, and the weather data included the location, altitude, and intensity (categorized into six levels) of a line-storm passage. The minimum distance between each aircraft and each level of precipitation was recorded every 6 s in the data file. This enabled the calculation of both the overall minimum distances to weather and also the accumulated durations in each level of weather for each aircraft.

As the pilots had access to on-board weather radar and also received reports of weather conditions by radio, the aircraft generally avoided the most severe regions of weather. The potential to translate this rerouting behavior into a form that could be incorporated into an automated decision aid was the motivation for this study.

4.2. Observed weather penetration

Of the 1095 aircraft, 353 (32%) penetrated level 2 weather or higher. Because the focus is on behavior of aircraft that penetrated weather, only the data for the 353 penetrating aircraft are considered here. Fig. 3 shows a cumulative distribution of the maximum amount of time that these 353 aircraft spent in levels 2–5. The solid lines show the observed duration values; the dashed lines show a model fit that is discussed in more detail below. Duration was defined as the accumulated time spent within a given level or a level of higher intensity. Time in level 2, for example, also includes time spent in levels 3 or higher, and thus serves as a metric of the total time spent within a region of precipitation.

Focusing on the solid lines, 90% of the aircraft, for example, spent less than approximately 190 s inside level 2, and 90% of the aircraft spent less than approximately 25 s in level 4. Additionally, the values of the cumulative distribution for duration of zero indicate the proportion of the 353 aircraft that did not enter each level of weather. For example, approximately 27% of the aircraft did not enter weather of level 3. No aircraft entered level 6.

The observed penetration times need to be corrected

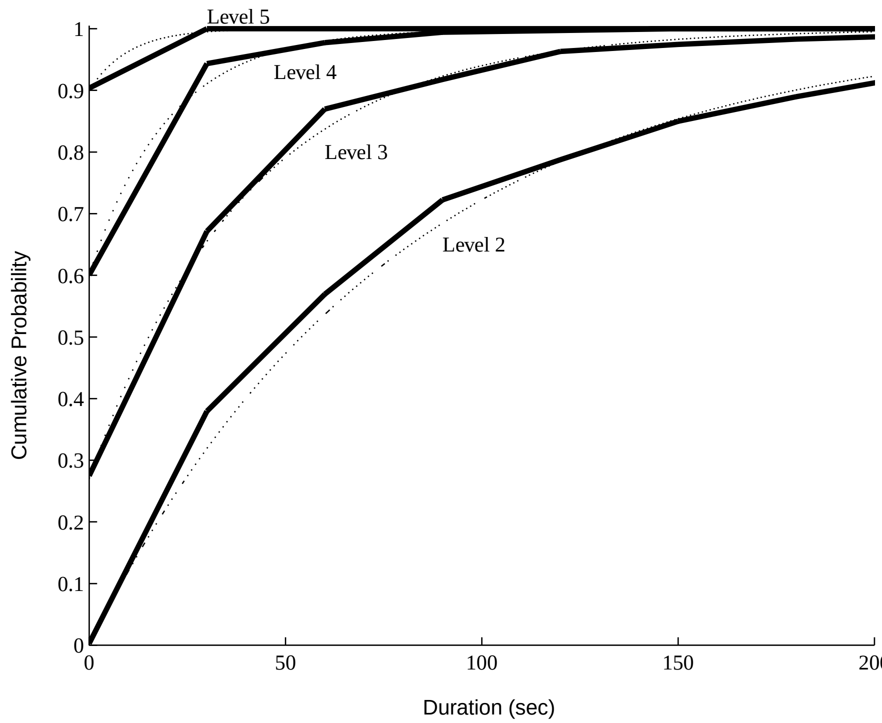

Fig. 3. Cumulative duration by precipitation level (solid lines: observed values; dashed lines: exponential hazard model fit).

to take into account the fact that the area covered by each level of precipitation varied. Level 5 weather, for example, covered only 9% of the area covered by level 2, and so shorter penetration times would be expected based solely on geometry. Modeling the weather as circular regions and assuming no deviation effects, the expected number of aircraft that would enter each level of weather would be proportional to the radius of each weather cell (or equivalently, the square-root of the area). Furthermore, the expected duration in each level, using this model, would be proportional to the area itself. Table 1 summarizes these relationships. The overall area covered by each level of precipitation is shown relative to the area covered by level 2. Also shown are the expected and observed fractions of aircraft that entered each level, and the overall expected and observed average duration in each level.

As can be seen in Table 1, increasingly fewer aircraft entered levels 4–6 than would be expected based on the simplified geometrical model of weather. Also, the average duration spent in levels 3 and above was lower than would be predicted by the model. Because on average the durations would have been significantly larger had no route modifications been made, the penetration times in Fig. 3 can be used as estimates of the upper limit of time that was acceptable to the pilots to spend in each level of weather.

4.3. Hazard modeling

Because none of the aircraft penetrated level 6 weather, level 6 can be adequately modeled as a hard hazard. Levels 2–5, however, had some degree of softness since aircraft did penetrate them. A simplifying assumption is that pilots penetrated the weather only as far as they considered to be acceptable. With this assumption, another way of interpreting Fig. 3 is that the cumulative distribution shows the probability that a pilot would *not* accept a routing of a given duration. Thus, 90% of the pilots, for example, would not

Table 1
Expected and observed penetration behavior (fractions relative to level 2 weather)

Precipitation level	Area fraction	Expected to enter	Observed to enter	Expected duration	Observed duration
2 (reference)	1.00	1.00	1.00	1.00	1.00
3	0.48	0.69	0.73	0.48	0.43
4	0.23	0.48	0.40	0.23	0.14
5	0.09	0.30	0.10	0.09	0.02
6	0.03	0.17	0.00	0.03	0.00

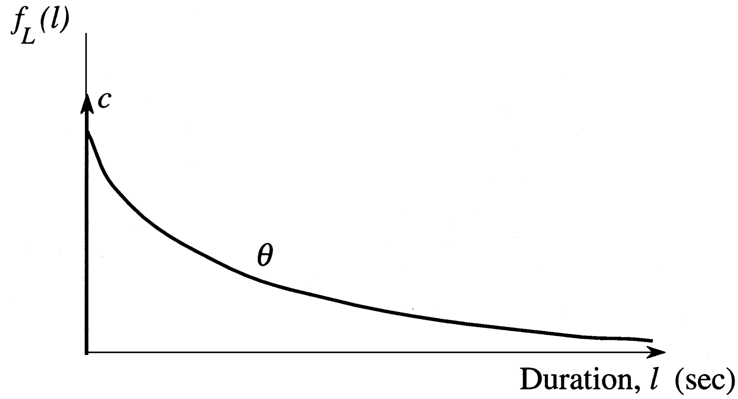

Fig. 4. Probability density function model for weather severity.

accept a trajectory that remained inside level 4 for 25 s. Similarly, since no aircraft were observed flying more than 150 s through level 4 precipitation, a trajectory that involves more than 150 s of flight through level 4 would not be acceptable to any pilot. This assumption is reasonable given the fact that the pilots, on average, originally had significantly longer trajectories through each level of precipitation, but deviated to reduce that exposure according to the cumulative plot in Fig. 3.

Using this approach, hazard severity PDFs, $f_L(l)$, for levels 2–5 were chosen such that the resulting probability of incident for a given length of exposure was similar to the observed cumulative distribution in Fig. 3. The result is that the computed probability of incident along a particular trajectory approximates the percentage of pilots who would not accept that trajectory. Thus, a decision threshold can be set based on a desired acceptance percentile. For example, assume that a decision threshold is set at a value of $P(\mathrm{I}|\mathrm{E})$ of 0.95. Then, alerts will be generated for 95th percentile weather; that is, weather for which 95% of pilots would agree is hazardous. A more risk-averse approach would be to lower this threshold, but then there may be a significant number of pilots who feel the system is overly conservative.

A slightly modified exponential hazard density function was used to model each precipitation level (Fig. 4). For each precipitation level, there is some discrete 'cost' or probability of incident that applies any time that type of weather is entered (modeled as an impulse of probability c at zero duration). In addition, there is an exposure time-dependent component that is integrated over the path of the weather, described by a mean time to an incident, θ. The result is the set of dashed lines in Fig. 3 which relate route acceptability (analogous to $P(\mathrm{I}|\mathrm{E})$) to the projected duration of the route in weather. Specific values of c and θ are shown in Table 2.

Table 2
Modeled hazard severity PDF parameters

Precipitation level	c	θ (s)
2	0.00	78
3	0.27	40
4	0.60	20
5	0.90	10
6	1.00	–

Due to the complex nature of weather as a soft hazard, the specific model of weather presented here is rudimentary and is intended only as an illustration of the type of analysis that could be pursued. Still, relatively accurate fits to the empirical data were possible using the time-dependent soft hazard model derived above. Future research efforts will focus on further developing this methodology and on analyzing aircraft–weather interactions in more detail, with the goal of enabling the development of decision-aids that integrate hard and soft hazard information in a manner that is acceptable to the operators.

5. Conclusion

This paper presents a generalized model that enables incorporating hard and soft hazards into a single automated decision regarding the acceptability of a particular state trajectory. The model may be used with more than one hazard, of more than one type, in a given problem. Uncertainties in state measurements; dynamics, hazard extent, and severity are included, as is consideration of the fact that different operators may have different concepts of what is an acceptable or unacceptable risk. This potential difference in the definition of what is acceptable is a key issue that needs to be resolved when developing decision aiding systems to monitor soft, subjective threats. In many cases, it is necessary to obtain data during the operation of the system in order to better understand operator preferences and decision-making behavior. This operational data can then be inserted into the design process of future decision-aids.

Using an aviation weather case study, information gained from observations of pilot behavior in the presence of weather was used to develop a preliminary model of weather as a soft hazard. This information could then be used in automation to aid operators in monitoring or replanning routes. Additionally, information such as traffic conflict probability (a hard hazard) can be combined directly with the soft hazard information to provide an overall assessment of the acceptability of a route.

Although an aviation case study was used here, the concepts developed in this paper can be extended to non-aerospace applications in which subjective operational data is inserted into the design process to improve automated decision-making. Examples include process control in which a soft operating envelope may be exceeded with the result that component failure rates may be increased. Whether an operator would agree to such an envelope exceedance could be determined and used with other dynamic information to develop a decision-aid.

The ability to manage soft or subjective hazards in automation could open up new design options to be considered. In automation today, hazards are generally defined to be 'hard' primarily due to either a clearly-defined physical limitation (e.g. collision), or due to the desire to enforce procedural limits on the operation of the process. Some of these hard limits may be imposed to reduce a complex soft hazard into a more easily managed form. If, however, soft limits can be defined and used in decision-aiding system algorithms, it may be possible to develop automation that is more flexible and more operationally acceptable when faced with such hazards.

Finally, the ability to manage soft hazard information may lead to new opportunities in developing decision-aids that are tailored to specific user preferences. Operators could specify their level of risk aversion or acceptance, which would then be used to determine when a soft hazard warrants alerts or other action by the automation. Alternatively, the soft hazard information could be used to provide feedback on the acceptability of a given route or operating strategy, both in terms of the overall risk level and by displaying the degree of softness or hardness of hazards that may be encountered.

Acknowledgements

This research was supported by the NASA Ames Research Center under Grant NAG-2-1111 and Contract NAS-2-98012. The authors are appreciative of the support provided by MIT Lincoln Laboratory in providing and analyzing the weather and traffic data, with specific thanks to James Evans, Dale Rhoda, Margo Pawlak, and Beth Bouchard.

References

[1] Davis T, Krzeczowski K, Bergh C. The final approach spacing tool. 13th IFAC Symposium on Automatic Control in Aerospace, Palo Alto, CA, September 1994.
[2] Slattery R, Green S. Conflict-free trajectory planning for air traffic control automation. NASA TM-108790, Moffett Field, CA, January 1994.
[3] Paielli RA, Erzberger H. Conflict probability estimation for free flight. AIAA J Guid Ctrl Dyn 1997;20(3):588–96.
[4] Krozel J, Weidner T, Hunter G. Terminal area guidance incorporating heavy weather, AIAA Guidance, Navigation, and Control Conference, August 1997.
[5] Kuchar JK. Methodology for alerting-system performance evaluation. AIAA J Guid Ctrl Dyn 1996;19(2).
[6] Pritchett A. Pilot non-conformance to alerting system commands during closely spaced parallel approaches. ScD Thesis. Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, December 1996.
[7] Rhoda DA, Pawlak ML. An assessment of thunderstorm penetrations and deviations by commercial aircraft in the terminal area. MIT Lincoln Laboratory Technical Report NASA/A-2, Lexington, MA, 1999.
[8] Kuchar JK, Carpenter BD. Airborne collision alerting logic for closely-spaced parallel approach. Air Traff Ctrl Q 1997;5(2).
[9] Yang LC, Kuchar JK. Prototype conflict alerting logic for free flight. AIAA J Guid Ctrl Dyn 1997;20(4).