Identifying and Modeling Safety Hazards

Jesse Daniels
Terry Bahill, Fellow of INCOSE
Systems and Industrial Engineering
University of Arizona
Tucson, AZ 85721-0020
danielsj@engr.arizona.edu
terry@sie.arizona.edu

Paul W. Werner
High Consequence Surety Engineering
Department 12331
Sandia National Laboratories
Albuquerque, NM 87185-0490
pwwerne@sandia.gov

ABSTRACT

The hazard model described in this paper is designed to accept data over the Internet from distributed databases. A hazard object template is used to ensure that all necessary descriptors are collected for each object. Three methods for combining the data are compared and contrasted. Three methods are used for handling the three types of interactions between the hazard objects.

INTRODUCTION

One of the major concerns of the Federal Aviation Administration (FAA) is creating and maintaining safe operating conditions for the airlines and the flying public. Recent events have prompted more rigorous analysis of hazards and their impact on airline safety. This paper describes a model that is intended to aid FAA inspectors in the detection and analysis of aviation-related hazards. However, the modeling technique is general and our results should have an intuitive and straightforward interpretation for users in most fields.

Currently, information on conditions present in the environment is gathered and stored in various databases. We intend to extract appropriate data from these databases and feed the data directly into the model for computation. An indirect goal of using such a model is determining which data are beneficial in assessing the impact of hazards on airline safety, i.e., the model can help decide where to allocate resources so that pertinent data are collected and frivolous data are rejected.

Another advantage in using such a model is that hazards can be quickly identified, described, categorized and archived for future reference. This provides a basis for which a consensus among experts can be attained. Thus, ambiguity associated with interpreting attributes of the hazards will be reduced, if not eliminated. In order to ease the archiving process we have defined a convenient format, called a *Hazard Object*, in which all necessary information can be captured. We will introduce the concept of hazard objects in a later section and explain how they are useful in evaluating hazardous conditions.

THE HAZARD MODEL

The structure of the Hazard Model is much like a traditional Systems Engineering tradeoff analysis (see http://www.sie.arizona.edu/sysengr/pinewood/ for an example). Related individual hazards are grouped into categories, whose cumulative relevance spans the space in which hazards may occur. Possible categories are Human Factors, Physiology, Weather, etc. Within each main category are the individual hazards that act as subfigures of merit in the tradeoff study analogy. The scores for each hazard are rolled up to form a score for their respective main category. Similarly, the scores for the main categories are combined to generate the overall score, called the *Hazard Index*, which functions as a numerical designation indicating how hazardous the situation is. A fragment of our model is shown in Table 1.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Table 1. Fragment of Sample Hazard Index Calculations

Figure of Merit							Input value	Score	Weight	Score times weight
1. Human Performance							0.34	0.76	0.50	0.38
2. Environment							0.87	0.96	0.50	0.48
2.1 **Weather**										
	Qualitative weight	Normalized weight	Input value	Scoring function	Score	Score times weight				
2.1.1 **Air Temperature**	10	0.50	120°F	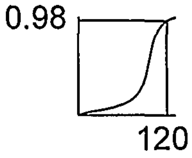	0.98	0.49				
2.1.2 **Visibility**	5	0.25	1 Mile	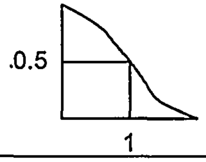	0.50	0.13				
2.1.3 **Relative Humidity**	5	0.25	100%	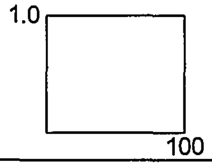	1.00	0.25				
Sum						0.87				
Hazard Index										0.86

This hierarchical partitioning of the hazards allows us to evaluate contributions of the individual hazards to the main category and each main category to the overall hazard assessment. The hierarchical structure also helps in dealing with dependence issues between hazards. Since hazards in a main category are intrinsically related, dependence between them remains isolated in that specific category and will not directly affect the overall evaluation.

HAZARD OBJECTS

Figure 1 shows an example of a hazard object that we will use to describe the details of hazard objects.

First, we give the hazard a descriptive *Title* so it can be easily and unmistakably identified. In this case, we are attempting to model the effects of "Pilot Recurrent Training" on safety. The next entry contains a *Description* of the hazard. This was done to ensure a consistent definition for each hazard. The *Units* of measure are included to indicate how the quantity is to be measured and how it is presented in graphical representations, such as scoring function charts. In this example, we chose to define the units of measure to be the number of hours of recurrent training in the current year.

The *scoring function* is given next as a chart on the hazard object as shown in Figure 1. Scoring functions (Wymore, 1993) are used to subjectively ascertain the severity of a hazard given the state of nature. Specific data on each identified hazard object is input into the scoring function and a score is given as output. Other facets of the Hazard Model then use these scores to generate an overall appraisal of the situation.

The scoring function chart is set up such that the input values (in the units of measure defined earlier) are displayed along the horizontal axis, and the output value (always between 0 and 1) is presented along the vertical axis. The function itself is defined so that the more hazardous the input value, the higher the output score will be. This is apparent from looking at the example since as the number of recurrent training hours in the current year diminishes, a higher score is output indicating a greater hazard.

A researcher will poll experts in the field where the hazard applies to get appropriate units of measure, reasonable ranges for input values, the scoring function shape and other characteristics of the scoring function. This captures the experts' perception concerning the hazard. There has been extensive research in developing a robust array of mathematical equations to represent scoring functions useful in such characterization. Examples of this will be given later.

The *Category* field initiates the hierarchical decomposition process. Categories are the largest subdivisions in the hazard space. It is here where we group hazards of similar characteristics such as environment-related hazards or human performance issues, etc. The Recurrent Training Hazard Object falls under the main category Human Performance. Next, we have the subcategory field that further breaks the categories down into finer granularity.

Hazard Object

Title: Pilot Recurrent Training

Description: The number of hours the pilot has spent in flight training in aircraft and approved simulators concerning specified maneuvers and procedures (e.g., CRM, wind shear, etc.)

Units: Hours in Current Year

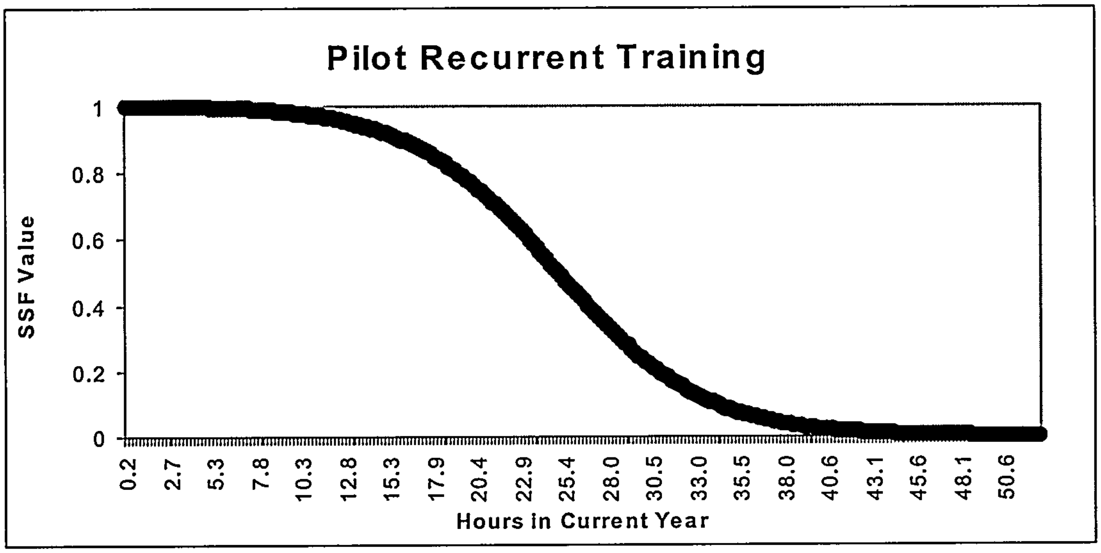

Category: Human Performance

Subcategory: Training

Phase of Flight: All

Interactions: N/A

Qualitative Weight: 6

Default Input Value: 5 hours in current year

FAR: 121.427

Aircraft Specific? Yes

Source of Data: Human Resources

Owner: Paul Werner

Date Approved: 3 Sept. 1999, Meeting, The University of Arizona

Notes: Regulatory minimum = 25 hours per year

Figure 1. An example of a hazard object.

For example, in the Human Performance category, we may include subcategories such as Psychology, Physiology, Communication etc. Recurrent training is part of the Training subcategory of Human Performance in the above example. We include a field called *Phase of Flight* because most hazards are most important during specific phases of flight, such as takeoff or taxiing. Thus, when a new hazard object is defined, the researcher must consider the interaction between the hazard and the phases of flight. In the example, we decided that all phases of flight are affected by training. Therefore, we entered "All" in the Phase of Flight field.

A more general field called *Interactions* allows the researcher the freedom to document all other interactive phenomena that transpire between hazards. An example of this would be the interaction between air temperature and relative humidity. This interaction exists because the relative humidity does not become a hazard unless (A) the air temperature is cold enough to (1) cause icing (either on aircraft or on the ground), (2) impact radio communications, (3) reduce radar and laser propagation, or (4) decrease visibility, or (B) the air temperature is hot enough to adversely affect aircraft performance.

Next, we include a *Qualitative Weight* field. The qualitative weight is a numerical entry given by experts indicating how much the expert believes the hazard effects safety. The qualitative weight is a value between 1 and 10, with a 1 meaning the hazard is minimally detrimental to system safety and a 10 denoting that if the hazard is present, it can severely impact system safety. In the example, we have assigned a qualitative weight of 6 for the hazard, indicating that the hazard rates 6 out of 10 in overall consequence on airline safety.

The *Default Input Value* field helps handle missing data. If there are inconclusive or missing data, or if good guesses for the inputs are not available, the Hazard Model will automatically revert to the default values. In most cases, we have set up the default values to be nearly the worst possible. This was done in part to motivate the acquisition of the necessary data, and also to promote a "better safe than sorry" safety philosophy. In the example, the default-input value was chosen to be 5 hours in the current year. This is a low value that generates an undesirable score for the hazard.

Links between hazards and the Federal Aviation Regulations (FARs) will be useful in providing a solid reference to the hazards and recognized mitigation procedures. Therefore, we have included a *FAR* field for the particular hazard. If there is a FAR relating to the hazard, it should be listed here. If there is no FAR pertaining to the hazard, then this might spur an investigation as to why and whether one should be created. This cross-referencing between the FARs and the hazards will reinforce the connection between regulations put in place to eliminate or mitigate hazards and the hazards themselves. It will also be useful in uncovering inadequacies in the FARs. Such is the case when a hazard is identified that is not attended to in the regulations.

The next field, *Aircraft Specific?* is a yes/no question used for indicating whether a hazard is specific to certain aircraft. Some hazards are common to all aircraft while some concern specific aircraft. This field will allow the researcher to illustrate such discrimination. In the example, it seems reasonable that pilot recurrent training may very well relate to the specific aircraft the pilot flies, and therefore will receive a "yes" in the "Aircraft Specific?" field.

The inputs for the scoring functions will be collected from numerous sources, so it is important to state where the data will come from. Therefore, we have provided a field, *Source of Data*. Examples of entries to this field could be databases, experts, reports, etc.

The owner of a hazard object is the person or group ultimately responsible for ensuring the accuracy and correctness of the information presented in a hazard object. The *Owner* field is provided so that if there are any discrepancies, questions or updates for information on the hazard object, the responsible party can be contacted. Typical entries here will include a name and telephone number.

The *Date Approved* field is self-explanatory. This field will also reference the source of the expertise: For example it will state whether the information in a hazard object came from a FAR, a book, one expert, two experts, a consensus of experts, a specific meeting, etc., along with the physical location when the decision was made. In this case, the Hazard Object was created during a meeting at The University of Arizona on September 3, 1999.

Finally, the *Notes* field is included to catch any information not addressed in the other fields. This may include explanations for information or special cases.

METHODS FOR COMBINING DATA

In this section, we explore three different methods for combining data to calculate the Hazard Index, which is a single numerical designation representing the state of system safety. The combining techniques described here are used at all levels of the system decomposition. At the lowest level when we are dealing with individual hazard objects, the scores are derived from the scoring function for each hazard object and the weights are based on expert opinion. These scores are used to comprise the hazard's subcategory score.

When we move to the next level and are dealing with main categories the scores are no longer directly derived from scoring functions, but are taken as the scores for each of the subcategories contained in the main category. The weights at this level may be evenly divided between the main categories unless there are obvious reasons to consider some main categories more significant than others. So again, the weights are based on expert opinion. Combining the data at this level yields the Hazard Index. The three data-combining methods are outlined next.

Linear Combination

The linear method of combining data is the simplest method of the three. It is useful when the output is not computed until the whole data set is available and when missing data are unlikely. The data combining process is as follows: Suppose there are *n* reasonably independent constituents to be combined. We assign a qualitative weight to each of the *n* constituents and then normalize the weights so they add up to 1. The score (valued from 0 to 1) associated with each constituent is then multiplied by the corresponding weight. The final result is the summation of the weight-times-score for each element. This process is commonly used, for example, when computing a grade point average for a student at a university. This technique is shown in Table 1.

The equation defining the process mathematically is given as

$$f = \sum_{i=1}^{n} w_i \cdot x_i ,$$

where *n* is the total number of elements to be combined, w_i represents the normalized weight and x_i represents the score for the i^{th} constituent. An extensive example of such rolling-up of figures of merit is given at http://www.sie.arizona.edu/sysengr/pinewood/.

Soft Aggregation

Arlin Cooper of Sandia National Labs developed the soft aggregation method discussed in this section (Cooper, 1999). The application of soft aggregation techniques is effective when incorporating uncertainty into the models and when the output is to be updated with each new input.

The soft aggregation model presented here attempts to combine information nonlinearly to obtain a single numeric output that characterizes the state of system attribute(s) (in this case the system attribute is safety). In particular, the aggregation is accomplished by utilizing an exponential combining function. The exponential technique was favored in part since the collective effects of combining such measures responds less and less to additional inputs, hence the name "soft aggregation".

The original model incorporates two families of information, those that lead to an increase in safety and those that lead to a decrease in safety. In this model, we only consider those that increase hazards, and thus decrement safety.

The function that serves our needs for a soft aggregation calculi is given as

$$f = 1 - e^{-\sum_{i=1}^{n} k w_i x_i}$$

The w_i indicate weights between 0 to 1 that suggest the significance of the measure with respect to increases in the overall hazard level for the situation. The x_i are "scores" affiliated with the i^{th} hazard. The k is a scaling constant used to further manipulate the results to match the requirements necessary for accurate evaluation. If k=1 the output of this model will range 0 to 0.63.

Certainty Factors

Certainty Factors (CFs) have been used successfully in the expert system arena for many years (Buchanan and Shortliffe, 1984). The underlying theory surrounding CFs is based on probability theory and has survived thorough mathematical scrutiny. Thus a vast knowledge base has developed for CFs, and a great deal is known about their properties and uses. As in the soft aggregation technique, the CF method is useful when the output is incrementally updated with each input and may also be used effectively when missing data or uncertain data cannot be avoided.

The initial CF data points are derived from the weights and scores for the hazard objects as follows: $CF_i = w_i * x_i$, where w_i is the normalized weight and x_i is the score (output of the scoring function) corresponding to the i^{th} hazard object. The CFs for the hazard objects are then combined to create a score for the associated subcategory. When we move to the next level up, the x_i's and w_i's become the weights and scores for the subcategories, and are combined to form a main category score.

An advantage to using CFs is that the weights do not have to be normalized. This means each time a new constituent is introduced, be it a new hazard object, subcategory, or main category, it is not necessary to re-normalize the weights. This artifact eases necessary computation consideration and resources.

The recursive formula for computing CFs is given as:

$$CF_{Aggregate} = CF_{Old} + (1 - CF_{Old}) * CF_{New},$$

where CF_{New} is the certainty factor for the new piece of evidence to be combined with the existing set and CF_{Old} represents the CF for the existing data set. $CF_{Aggregate}$ is the cumulative outcome from combining the data. At the highest level, this becomes the Hazard Index. In our system, the CFs are restricted to the range [0, 1].

In this project we evaluate two-dozen case studies with these three techniques and we show the differences and similarities.

INDEPENDENCE & INTERACTIONS

As an example of independence and interactions, consider modeling the hunting behavior of lions in Africa. First, consider independence. It would be incorrect to use both weight and sex of the lions to predict success, because these measures are dependent.

Next, consider interactions. If one lion chases a Thompkins Gazelle, it has a one in eight likelihood of catching the gazelle. If another lion chases another gazelle, it has a one in eight likelihood of catching it. We expect a *linear addition* for the two hunting sprints, which gives a one in four likelihood of success. Now consider two lions simultaneously chasing one gazelle. They have a one in two likelihood of catching it. *Cooperation* gives a greater likelihood of success than linear addition. Next, suppose a cheetah is chasing a gazelle and a lion joins in the chase. The likelihood of success becomes about one in twelve. *Interference* decreases the likelihood of success below that of linear addition.

In terms of hazards when the effect of two hazards is greater than that obtained by linear addition, such as mixing Clorox and Sani Flush producing chorine gas, we call the effect an *amplifying interaction.* When the effect of two hazards is less than that obtained by linear addition, such as the presence of flammable liquids and a temperature below freezing, we call the effect an *attenuating interaction.*

When selecting components for the hazard model, the components should be as independent as possible. It is more difficult to analyze systems when the metrics depend on one another. However, when analyzing complex systems this is almost unavoidable.

In the Hazard Model, we have dealt with the interaction issue in two ways. For attenuating interactions, the final score for the first measure is the product of the original score for the first measure and the score for the second. This technique allows us to model phenomena where if one constituent of the interactive relation is close to zero, then the hazard is close to zero. A good example is the relative humidity – air temperature interaction. If the temperature is below freezing, then high humidity increases the chance of icing and therefore increases the hazard. However, at 70°F high humidity is no hazard at all.

Amplifying interactions involves metrics in which there is a synergistic effect. That is, when two metrics are present the hazard from the combination is greater than the hazard resulting from the linear sum of the two metrics. For such amplifying interactions we use an equation from reliability theory: $H_{Overall} = 1-[(1-H_1)(1-H_2)]$. In this equation, $H_{Overall}$ represents the resulting score after the interaction effect has been accounted for. H_1 and H_2 are the original scores for the first and second hazards, respectively. This equation gives us the synergistic effect desired. An example of an amplifying interaction is consumption of alcohol and barbiturates. The combined effect of the alcohol and barbiturates is greater than the effects modeled by the linear sum. Another example is chopsticks: a pair of chopsticks performs much better than the linear sum of two individual chopsticks.

CONCLUSIONS

The hazard object template presented in this paper would be a good template to use in any database. The field for *Units* is particularly important: it may seem obvious, but experience has shown that it is easy to omit. We showed three methods for combining the data. In the final paper, we will have experimental results comparing these three techniques. It is hard for humans to deal with interactions between objects. But it is easy for computers to do so. We have presented three types of interactions: linear addition, amplifying and attenuating. We gave examples of each and presented a mathematical treatment of each.

Acknowledgements. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

The authors gratefully acknowledge the support of the Federal Aviation Administration (AAR-424). This paper is based on work done at Sandia National Laboratories and University of Arizona.

REFERENCES

Cooper, J. A., "Soft mathematical aggregation in safety assessment and decision analysis," *Proceedings of the 17th International System Safety Conference,* August 16-21, 1999, Orlando, FL.

Buchanan, B. G. and Shortliffe, E.H., (Eds.) *Rule-Based Expert Systems*, Addison-Wesley, Reading, MA, 1984.

Wymore, A. W., *Model-Based Systems Engineering*, CRC Press, Boca Raton FL, 1993.

AUTHOR BIOGRAPHIES

Jesse Daniels is a graduate student at the University of Arizona in the Systems and Industrial Engineering Department. He received his B.S. in systems engineering in 1999 from the University of Arizona.

A. Terry Bahill is a Professor of Systems Engineering at the University of Arizona in Tucson. He received his Ph.D. in electrical engineering and computer science from the University of California, Berkeley, in 1975. He holds a U.S. patent for the Bat Chooser™, a system that computes the Ideal Bat Weight™ for individual baseball and softball batters. He is Editor of the CRC Press Series on Systems Engineering. He is a Fellow of the Institute of Electrical and Electronic Engineers (IEEE) and of the INCOSE. He is the chair of the INCOSE Fellows Selection Committee.

Paul W. Werner is a Principal Member of the technical staff at Sandia National Laboratories and works in the High Consequence Surety Engineering Department. He received his Ph.D. in electrical engineering from the University of New Mexico in 1991. He is also a Commander and a Naval Flight Officer in the U.S. Naval Reserve.