UNCERTAINTY IMPORTANCE ANALYSIS FOR AVIATION EVENT TREES

John Shortle, Seungwon Noh, George Mason University, Fairfax, VA

Abstract

Event trees are often used in the assessment of aviation safety. Since catastrophic accidents are rare, the number of observed events to quantify the model may be small or even zero. In particular, some sequences of an event tree may have never occurred. This has the potential to create large uncertainty ranges for the parameters estimated. The objective of this paper is to identify the most critical nodes in an event tree with respect to their potential for reducing the overall uncertainty of the model. To do this, an uncertainty important metric from the literature is applied. The metric answers the question: If one were able to know the exact value of a parameter, say, through a large data collection effort, how much would the overall uncertainty of the model improve and which parameter would lead to the best uncertainty improvement? We provide several case studies of the method to existing event sequence diagrams in the Integrated Safety Assessment Model. Finally, we consider the method applied to hypothetical event trees with no observed events, such as those that might apply to future airspace concepts. Results show that nodes that are closest in proximity (within the model structure) to a catastrophic end event tend to be the most critical. One application of these results is to prioritize nodes for quantification via expert elicitation exercises.

Introduction

Event trees are often used in the assessment of aviation safety. Advantages of event trees are that they are straight-forward to evaluate and easy to understand [1]. One challenge is that there may be little historical data that can be used to quantify the parameters of the model. Since catastrophic accidents are rare, the number of observed events available to quantify the model may be small or even zero. In particular, some sequences of an event tree may have never occurred, and thus appropriate probabilities along those event paths are estimated to be zero.

However, just because an event has not occurred in the past, does not mean that it will never occur in the future. This means that there is some uncertainty around the probability estimates that are quantified via data. This issue is particularly significant for rare-event problems, because the uncertainty ranges can be very large relative to the true underlying probabilities. Even though an uncertainty range can be reduced with more observations – for example, by observing the system over a longer time horizon – there is always some degree of uncertainty, regardless of how many observations are collected.

This paper applies an uncertainty importance metric [2] to event-tree models in aviation. The goal is to determine which parameters are most critical to reducing the overall uncertainty of the model. That is, if one were able to obtain more observations related to a specific parameter, how much would the overall uncertainty of the model improve and which parameter would lead to the best uncertainty improvement?

We provide several case studies associated with event-sequence diagrams in the Integrated Safety Assessment Model (ISAM), developed by the FAA [3]. We also apply the uncertainty importance methods to event sequence diagrams with no observed events. For example, event trees for future airspace concepts may have little or no current data. We demonstrate that if a rough order of magnitude for some of the end events can be estimated, then it is often possible to identify the most critical nodes of the event trees under very general conditions even with no historical data.

One application of these results is to prioritize nodes for quantification via expert elicitation exercises. Such exercises are expensive to conduct and yield results only for a limited number of nodes. Thus, identifying the nodes that would yield the greatest reduction in overall model uncertainty may be helpful.

Uncertainty in Event Trees

This section describes how uncertainty can be quantified for parameters in an event tree. Parameter uncertainty depends on the sources of data to quantify the model.

As an example, consider the event tree associated with a mid-air collision in Figure 1. This is a *notional* event tree meant to illustrate basic concepts. The initiating event is a situation in which two aircraft have lost separation (that is, they are laterally within 5 nautical miles of each other and vertically within 1,000 feet of each other), *and* they are on a collision course. The first safety layer is air traffic control (ATC) which can provide resolution maneuvers to the flight crew in order to avoid any close proximity event. If this action is successful, there is a loss of separation (LOS), but otherwise the flights continue without any further consequences. If this action is not successful, then the next safety layer is a resolution by the flight crew via TCAS (Traffic Collision Avoidance System). If the flight crew is successful in resolving the conflict, then it is assumed that the aircraft miss each other, but there is a near mid-air collision (NMAC, defined as two aircraft within 500 feet laterally and 100 feet vertically). If the action is not successful, then there is a collision. This is a notional event tree; there are many embedded assumptions, and the tree avoids complicating factors such as encounter geometry and the precise timing of these events. (For an example of an event tree that considers the timing of events, see [4].)

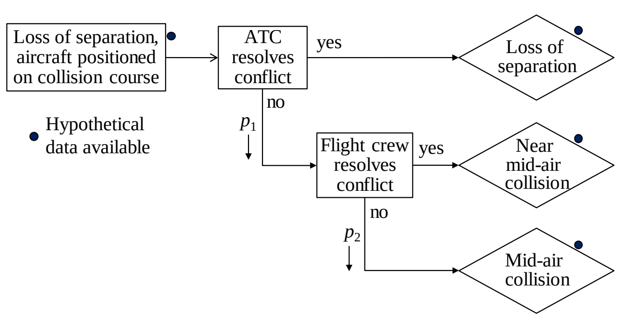

Figure 1. Event Tree for Mid-air Collision

For the purposes of data collection, suppose that data are collected for the initiating event and the three end events (loss of separation, near mid-air collision, and mid-air collision). Because every path through the event tree must terminate at one end event, the total number of end events must sum to the number of initiating events. That is, if we know three of the four numbers, then we can derive the fourth.

For example, suppose that the following *hypothetical* event counts have been observed:

- Initiating events: 1,000
- Losses of separation: 990
- Near mid-air collisions: 9
- Mid-air collisions: 1

Naturally, most events terminate in the least severe outcome, a loss of separation. Fewer events terminate in a near-mid-air collision, and even fewer terminate in a mid-air collision.

In the figure, the probabilities p_1 and p_2 are the *failure* probabilities of the respective safety layers, which are sometimes referred to as *barriers* [5]. Based on this notional data set, the intermediate transition probabilities can be inferred from the event data. In this case, $p_1 = 0.01$ and $p_2 = 0.1$. These probabilities are consistent with the initiating event and end event counts.

As a side note, the observation counts are often reported as *frequencies* (e.g., events per operation, or events per flight hour). Event counts can be converted to frequencies by dividing by the total number of operations or total number of flight hours. For example, if these event data come from a set of 10^8 operations, then the initiating event occurs with a frequency of 1 per 10^5 operations, and mid-air collisions occur with a frequency of 1 per 10^8 operations. Regardless of whether the end events are quantified as counts or frequencies, the intermediate transition values p_1 and p_2 are unit-less probabilities with values between 0 and 1.

The key point is that the process of computing the intermediate probabilities p_1 and p_2 from the event data results in *point estimates* for the parameters. In reality, there is some uncertainty in these values. This is because there is randomness in the underlying processes that produce the original event data from which the probabilities are derived.

The uncertainty ranges are particularly large (in a relative sense) for rare events with small event counts. For example, if one were to observe another 100 million flights, one might just as easily observe zero mid-air collisions or two or three. There is quite a bit of variability that occurs by chance. In fact, if this hypothetical experiment were carried out many times, it would be odd to observe *exactly* one collision every time.

In many practical examples, the counts for many end events are zero. However, just because there are no observed events does not mean that such events cannot occur. In fact, there is a rule-of-thumb, called the "rule of three" which states that if no accidents have occurred over n operations, then the upper bound for the probability of an accident is about $3/n$ (more specifically, this is the upper bound for a 95% confidence interval on the probability). For example, if no accidents have been observed over 100 million (10^8) operations, then the upper bound on the frequency of an accident (per operation) is about 3×10^{-8} (or 3 per 100 million operations).

Zare and Shortle [6] have given a method for estimating parameter uncertainties in an event tree from the event counts. We summarize this method here. The basic steps are:

A. Obtain an uncertainty distribution for each parameter estimated via data (in this example, these parameters would be the number of initiating events and the numbers of end events). This process is described in more detail below.

B. Repeat the following Monte-Carlo simulation step:

   a. Generate a random event count for each end event based on the uncertainty distributions established in Step A. (In this example, during one execution of the Monte-Carlo step, we might get something like 1,011 LOS events, 8 NMACs, and 2 mid-air collisions. The total number of initiating events, by implication, would be 1,021.)

   b. From the event count data, compute the inferred intermediate probabilities. (Continuing the example, we get $p_1 = 10/1021 \approx .0098$ and $p_2 = 2/10 = .2$.)

   c. Tabulate the sample values of the event counts and intermediate probabilities.

In step A, the uncertainty distributions are derived based on the assumption that each event count is a Poisson random variable. (This is a common assumption for rare events. For events with large counts, a Poisson distribution is approximately normal.) See [6] for more details.

Figure 2 shows the resulting uncertainty distributions for the example event tree. As might be expected, the uncertainty ranges for the initiating event and the loss-of-separation end event are the narrowest, because these events are observed more often. The frequencies are known with a very tight confidence interval. In contrast, there is more relative uncertainty in the NMAC counts and even greater uncertainty in the mid-air collision count.

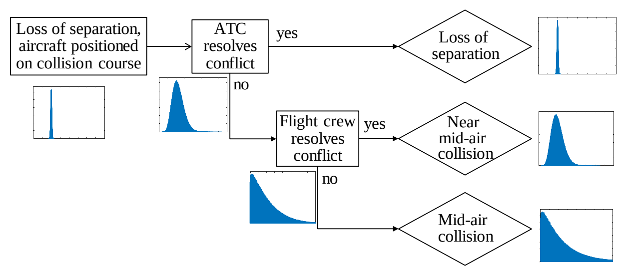

Figure 2. Parameter Uncertainty for Event Tree

To be more specific, when we say that an uncertainty distribution has a "narrow" confidence interval, we mean this in a *relative* sense – that the width of the distribution is small *relative to the mean of the distribution*. In Figure 2, the *x*-axis of each distribution is scaled to go from zero to four times the mean of the distribution. So each graph shows the same relative scale.

For example, the mean of the initiating event distribution is 1,000, which is the number of observed events. Based on the assumption of a Poisson distribution for the event count, a 95% confidence interval goes from about 940 to 1,060. In absolute numbers, the interval has a width of 120. In a relative sense, the width of the confidence interval is about 12% of the mean. In contrast, the mean of the collision count distribution has a mean of 1 and a 95% confidence interval that goes from about 0.05 to 4.74. This is a smaller range in absolute value, but it is much larger relative to the size of the parameter to be estimated. So we say that there is more uncertainty associated with the collision count than with the initiating event count.

Uncertainty Importance Metric

This section describes an *uncertainty importance* metric that has been used in the literature [2]. The metric quantifies, in some sense, how critical the uncertainty of a given parameter is and how beneficial it would be to reduce that uncertainty.

The metric answers the following question: If a given parameter X were known with exact certainty, how much would the uncertainty in the overall model output decrease? Said another way, if an analyst had a crystal ball that could provide the precise value for one (and only one) parameter in the model, which parameter value should the analyst ask of the crystal ball?

Definition: The *uncertainty importance* of a parameter X is the expected percent reduction in the variance of overall accident risk due to ascertaining the value of X with certainty [2].

Figure 3 illustrates the basic idea of the uncertainty importance metric applied to the parameter p_2 in the notional model. The parameter quantifies the failure probability of the flight crew resolving the conflict. As discussed previously, there is uncertainty associated with the estimated value of p_2. This uncertainty is characterized by a distribution. The idea is to replace this distribution with an exact value – in particular, the mean of the distribution. In Figure 3, this is represented by a point mass, or a probability spike, at this precise value.

Now, once p_2 is known exactly, then the associated uncertainty distributions in the downstream end events become narrower. In particular, the uncertainty associated with the probability of a mid-air collision, which is the critical output of the model, is reduced. The percent reduction in the variance of this distribution is the value of the uncertainty metric.

This process can then be repeated for other parameters of the model, such as p_1. Then whichever parameter has the largest uncertainty importance metric is deemed to be most critical in terms of reducing model output uncertainty.

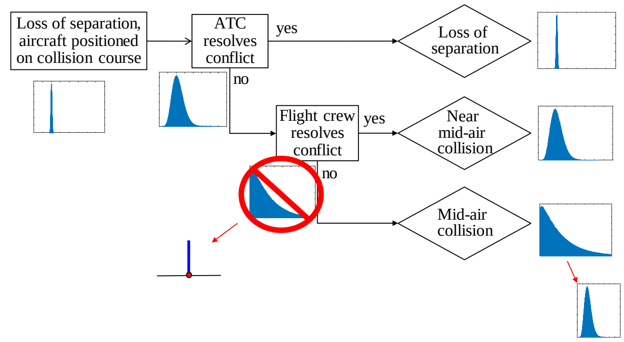

Figure 3. Uncertainty Importance Metric

In some sense, the uncertainty importance metric captures two basic ideas: Uncertainty and sensitivity. Uncertainty quantifies the unknown range of the parameter value. In general, there is more value in precisely quantifying a parameter with high uncertainty compared to one with low uncertainty, since the parameter with low uncertainty is already known fairly precisely.

Sensitivity quantifies how much the model output changes in response to changes in an input parameter. Parameters with high sensitivity, if they are changed by a little bit, have a large effect on the model output. Parameters with low sensitivity, if they are changed significantly, may have little to no effect on model output. In general, there is more value in precisely quantifying a parameter with high sensitivity, since small errors in the quantification can result in large model output errors. Noh and Shortle [7], [8] provide a discussion of various sensitivity metrics in the literature and provide an analysis of these metrics on the Integrated Safety Assessment Model, including an analysis of common cause failures.

In general, parameters with both high uncertainty and high sensitivity are the best candidates for further quantification efforts. The uncertainty importance metric essentially captures both ideas together. A parameter with high sensitivity tends to translate uncertainty in the parameter to uncertainty in model output, because small changes in the parameter translate to large changes in model output. Thus, eliminating uncertainty in the parameter has a large impact on reducing uncertainty in the model output.

As another example, a parameter with high uncertainty and low sensitivity might have a low uncertainty importance metric. Even though there is a large uncertainty in the parameter, it may not really matter, because the parameter has little impact on the overall model output, so any value within the uncertainty range gives about the same output.

Case Study

This section considers several event trees from the Integrated Safety Assessment Model (ISAM) as case studies for application of the uncertainty importance metric. ISAM models accident and incident scenarios in the National Airspace System through a set of event-sequence diagrams (ESDs) and supporting fault trees [2]. ISAM contains 35 ESDs, each corresponding to a different initiating event. In this section, we apply the uncertainty importance metric to three different ESDs – fire onboard aircraft, single-engine failure in flight, and unstable approach. (The analysis in this paper is applied to ISAM version 3.04; a newer version ISAM 4.0 also exists. The textual description of some events have been modified slightly from the model for clarity.)

Figure 4 shows the event sequence diagram for fire on board aircraft. In this tree, there is one extreme outcome (end event 1) – collision with the ground. The other three outcomes are less severe – personal injury, aircraft damaged but continues flight, and aircraft continues flight. For the sake of example, we say that end event 1 is a "severe accident" while the other end events are less severe. The goal is to evaluate the model with respect to estimating the probability of a severe accident.

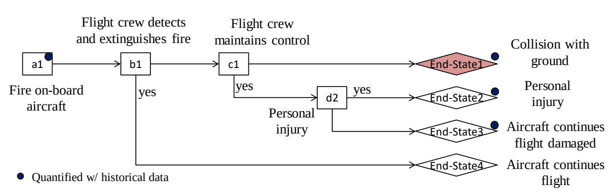

Figure 4. Event Tree for On-board Fire

In this tree, event counts exist for all four end events (the count of the fourth end event is the difference between the number of initiating events and the total of the other three events). Thus all intermediate parameters of the model are uniquely determined from the observed event counts.

Table 1 shows the resulting uncertainty importance metrics. The absolute numbers are not critical; it is the relative ranking that is more important.

Table 1. On-board Fire Uncertainty Importance

Parameter	Uncertainty Importance	Uncertainty (CV)
c1	.951	.665
b1	.070	.152
d2	.000	.166
a1	.000	.007

As might be expected, parameter c1 ranks the highest. This is because the outcome of this node most directly determines whether or not there is a severe accident. If the event path branches to the right (flight crew does not maintain control), there is a catastrophic accident; if the event path branches down, then all possible downstream end events are classified as less severe.

The last column of the table shows the uncertainty of each parameter defined as the standard deviation of the parameter estimate divided by its mean, also known as the coefficient of variation (CV). Parameter c1 also happens to have the largest uncertainty. So the parameter scores high on two counts: (1) The parameter has a big impact on whether or not a severe accident occurs, and (2) the parameter has a large relative uncertainty. Thus, if we knew its value exactly, this would go a long way in terms of reducing uncertainty associated with the model's estimate of the probability of a severe accident.

Parameter b1 is the next most important parameter. This node is important because it also differentiates between severe accidents and less severe events. However, unlike c1, a right branch does not always lead to a severe accident, so it does not directly distinguish between the two cases.

Parameter d2 has an uncertainty importance metric of zero. By definition, since the uncertainty importance metric relates only to the probability of a severe accident, and since both outcomes of d2 are less severe, it provides no impact on the probability of a severe accident. Parameter a1 has a low uncertainty importance metric because its value is known fairly precisely to begin with.

Now we consider a slightly more complicated example. Figure 5 shows the event sequence diagram for single-engine failure in flight. In this tree, there are three severe outcomes (end events 1, 2, and 4) – namely, collision with the ground following total loss of engine power, landing off runway (i.e., unable to reach any nearby airport), and collision with the ground following partial loss of engine power. The other two outcomes (end events 3 and 5) are much less extreme – landing safely at an alternate airport (but unable to reach the original destination) and continuing flight (i.e., landing safely at the original destination).

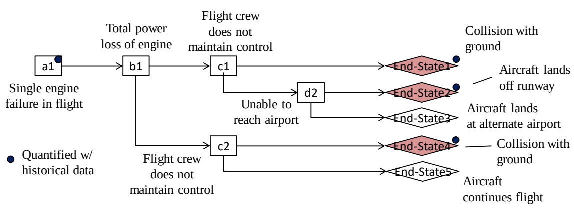

Figure 5. Single-engine Failure in Flight

The methods described previously have additional complications when applied to this event tree. The methods assume that event counts are available for every end event. In practice, event counts may only be available for moderate-to-extreme end events. Data on low consequence events may not be available.

Figure 6 shows the nodes with available event counts. The extreme events are all recorded. And the initiating events are recorded (i.e., whenever there is an engine failure). But whether or not an aircraft safely lands at its original airport or diverts to a nearby airport is not in the data set. In other words, we can infer the sum of end states 3 and 5, but we do not know how this total is divided between the two outcomes. Thus, the parameters in the model are under-determined.

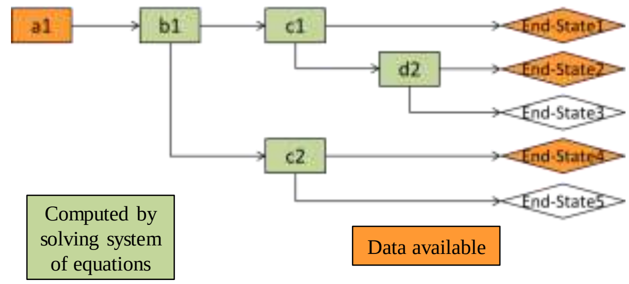

Figure 6. Event Tree with Incomplete Data

To handle this situation, we make the following adjustment to the previously described algorithm: In the Monte-Carlo step, we first generate a random event count for the initiating event and each end event which is supported by data. From this, the total event count for the remaining two end events (end events 3 and 5) can be inferred. The event count for one of the events is determined by a random draw from a uniform distribution over this total amount, and the remaining count is assigned to the other event. Once all of the end event counts are obtained, the intermediate probabilities (b1, c1, c2, d2) can be inferred by solving a system of equations.

The resulting uncertainty importance metrics are given in Table 2. As might be expected, the most critical parameters are c2 and d2, which are the last branches in the tree that differentiate between severe accidents and other events. For example, d2 determines whether or not an aircraft is able to make it to a nearby runway or not. Parameters that are earlier in the accident sequence, such as b1, are less important because, regardless of which path an event takes leaving b1, the system can still end up in either a severe accident or non-accident.

Table 2. Engine Failure Uncertainty Importance

Parameter	Uncertainty Importance	Uncertainty (CV)
d2	.~~5420~~6220	6.779
c2	.3168~~2856~~	9.074
c1	.0593~~1388~~	7.869
b1	.2487~~0390~~	0.577
a1	.0005~~7~~	0.009

The last column of the table shows the uncertainty of each parameter, defined as its coefficient of variation (CV). For example, c2 has a

CV of about 9, meaning that the standard deviation of the parameter estimate is nine times greater than the mean itself. Parameters with higher uncertainty are more likely to have higher uncertainty importance, since the importance metrics captures the effect of reducing that uncertainty to zero.

Figure 7 shows the event sequence diagram for an unstable approach. (This tree is from ISAM 3.04; the corresponding tree in ISAM 4.0 is slightly different and has more end events.) This tree has many more branches than the previous examples. The first branch (b1) corresponds to whether or not the flight crew initiates a rejected approach. Because of the complexity of the tree, we do not give a textual description for every node. The red end events are the most severe, including outcomes such as collision with the ground, undershoot of the runway, overshoot, and veer off. The non-shaded end events are less severe and include outcomes such as aircraft continues landing roll and aircraft continues rejected approach.

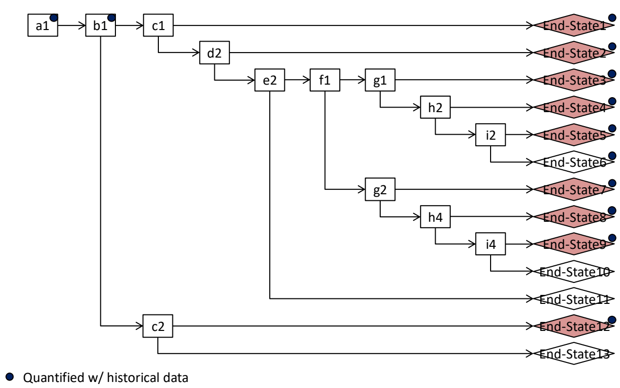

Figure 7. Unstable Approach

In the figure, the dots represent nodes with event-count data (some event counts may be zero). In addition to selected end events, nodes a1 and b1 are also quantified. That is, event counts exist for the number of unstable approaches (a1) as well as the fraction of these approaches that are rejected (b1). Knowing a1 and b1 gives the number of events reaching nodes c1 and c2. This means that the last end event (end event 13) is also quantified (equal to the number of events reaching c2 minus the event count for end event 12). There are two un-determined parameters – end events 10 and 11. As in the previous example, in the Monte-Carlo step of the algorithm, the event count for one of the events is determined by a random draw from a uniform distribution over the total between them.

Table 3 shows the uncertainty importance metric for this event tree. The most important parameters turn out to be h4, i4, and f1. These parameters are, respectively, insufficient runway length remaining, sufficient braking not accomplished, and structural failure. All three nodes are important because they have large uncertainties and differentiate between severe accidents (e.g., runway overrun) and less severe events (e.g., aircraft continues landing roll).

Table 3. Unstable Approach Results

Parameter	Uncertainty Importance	Uncertainty (CV)
h4	.726	94.22
i4	.187	61.24
f1	.087	61.28
g1	.022	0.61
e2	.005	0.58
i2	.001	3.02
h2	.0004	3.01
g2	.0001	70.95
b1	.0001	0.0003
a1	.00009	0.0002
c1	.00001	0.33
d2	.00001	3.14
c2	.00001	0.95

Interestingly, there are some parameters with large uncertainty importance, but low uncertainty, and vice-versa. For example, g1 ranks fourth for uncertainty importance, but has a relatively low uncertainty. The parameter g1 (flight crew des not maintain control) turns out to be important because it plays a critical role in differentiating between accidents and non-accidents. This is because there are no observations for end events 4 and 5, so g1, based on the data, is effectively the last branch to decide between a severe accident and less severe event. Even with relatively low uncertainty, there is still value in the model to know its value is exactly. Conversely, g2 (flight crew does not maintain control) has high uncertainty, but low uncertainty importance. The uncertainty is high because end state 10 is one of the un-determined parameters, so it can vary substantially. But since there are no observations for end event 7, the downstream nodes, h4 and i4, are more critical in differentiating between severe accidents and less severe events.

Models with Little Data

Event trees associated with future airspace concepts may have little or no historical data, since the concepts have not yet been implemented. Similarly, models associated with unmanned air vehicles may have less historical data, since there have been fewer years of data collection. To use the methods described in this paper, *some* starting data is needed to quantify the models.

In this section, we conduct a parameter sweep of the model parameters to identify general conditions under which a given parameter in the model is most significant. If a rough order of magnitude for some of the end events can be estimated, then it is often possible to identify the most critical nodes of the event trees under very general conditions even with no historical data.

As seen from the examples in the previous section, event trees can take on a variety of structural forms. In order to obtain results that can generalize to a large number of event trees, we consider the transformation of these event trees to an equivalent *barrier* representation. A barrier is anything that lowers the likelihood or consequence of an undesired state.

For example, the notional event tree from Figure 1, for a mid-air collision, is equivalent to the barrier representation model in Figure 8. Air traffic control (ATC) can resolve the conflict or the flight crew can resolve the conflict. If both barriers fail, then a mid-air collision occurs (starting from a state where two aircraft are positioned on a collision course with a loss of separation). The probabilities p_1 and p_2 represent the failure probabilities of each barrier. The probability that both barriers fail is $p_1 \cdot p_2$.

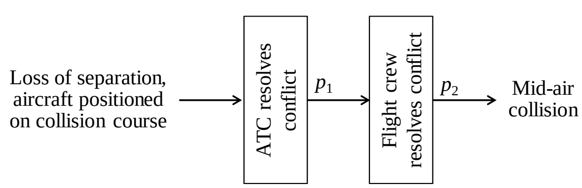

Figure 8. Barrier Representation of Event Tree

In general, event trees associated with aviation safety can be quite complicated. In the Integrated Safety Assessment Model (ISAM), for example, the models are further complicated by adding a supporting fault tree beneath every node of an event tree.

Shortle and Noh [9] provide a general algorithm and heuristic to convert an event tree with underlying fault trees to a barrier representation model. The barrier model is mathematically equivalent to the original model, but sometimes requires "pruning" lower-level nodes in the fault trees in order to achieve this equivalence. The barrier model offers a higher level simplified version of the model. This method is applied to identify barriers in the event sequence diagrams from ISAM [9].

Figure 9 shows a notional diagram for a system with three barriers. Starting from the initiating event, the accident occurs if all three barriers fail. If any of the three barriers is successful, the accident is averted.

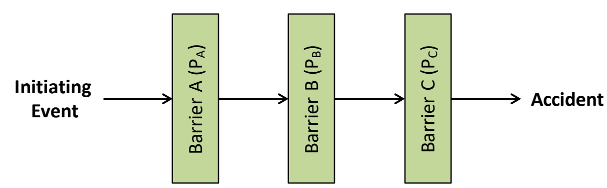

Figure 9. Barrier Representation of Event Tree

This model has four parameters to be quantified – the frequency of initiating events and the failure probability of each barrier, p_a, p_b, and p_c. The uncertainty importance of each parameter depends on the event count data. In this notional example, it is assumed that event counts exist for the initiating events, accidents, and success occurrences of each barrier. Given an assumed set of data counts, the uncertainty importance metrics for each parameter can be determined using the methods described previously. These metrics are different depending on what event counts are assumed.

To identify general conditions under which, say, parameter p_a is the most important, we conduct a parameter sweep for the event counts. Table 4 describes the parameter sweep for these experiments. For each combination of parameters, the uncertainty importance metric for each parameter is obtained and the most important parameter is determined. The total number of experiments is 4×12×12×12 = 6,912. For example, in one of the experiments, the number of initiating events is 10^4, $p_a = 0.1$, $p_b = 0.01$, and $p_c = 0.1$. The number of successes for barrier A is set to 9,000; the number of successes for barrier B is set to 990; the number of successes for barrier C is set to 9, and the number of accidents is set to 1. Based on this event count data, the uncertainty importance metrics are computed as before, without assuming any knowledge of the true values for p_a, p_b, and p_c.

Table 4. Experimental Setup

Variable	Values
Total flights	10^9
Initiating events	10^2, 10^3, 10^4, 10^5
Barrier failure probability	10^{-4}, 10^{-3}, 10^{-2}, .1, .2, .3, .4, .5, .6, .7, .8, .9

Figure 10 shows the results of this parameter sweep. The three axes correspond to the true failure probabilities of each parameter p_a, p_b, and p_c. For each parameter combination, the most important parameter is shown by the color coding of the respective point in the graph. There are three distinct zones corresponding to where a given parameter is most important.

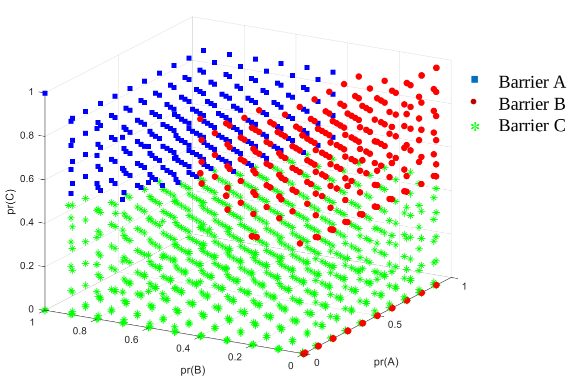

Figure 10. Parameter Sweep for Uncertainty

Of particular note, p_c has the largest uncertainty importance as long as its value is less than about 0.5. In other words, even if there is no data to populate the model, if one believes that the true failure probability of p_c is less than 0.5, then p_c is the most important parameter, regardless of the value of the other parameters. As an example, consider a model of a mid-air collision involving UAVs. The last barrier might be whether or not a maneuver to avoid an imminent collision is successful. The first barrier might be strategic conflict avoidance by the UAV operator. If we believe the system to avoid an imminent collision has at least a 50% probability of success, then this probability is the most critical model parameter to estimate precisely.

Since the barrier structure of Figure 9 is quite general, the results show that nodes or barriers that are closest in proximity (within the model structure) to a catastrophic end event are most critical. We expect that similar results would hold for models with more than three barriers.

Conclusions

Estimates for parameters in an event tree are often given as *point estimates*. Since rare-event estimates are inherently noisy, in the interest of model transparency, we believe that it is helpful to report *uncertainty intervals* associated with each estimate.

This paper described methods to quantify uncertainty in an event tree and to ascertain the potential reduction in model uncertainty from knowing the exact value of a given parameter. (The methods are adapted from similar concepts and methods described in the literature; Azin and Shortle, 2017; Iman, 1987). This provides some indication of which parameters are most critical for further efforts in either data collection, modeling, or expert elicitation.

The methods were applied to several case studies in the Integrated Safety Assessment Model. In addition, the methods were also applied to hypothetical event trees with no baseline data. Results tended to indicate that nodes that are closest to the catastrophic events, within the structural framework of the event tree, are the most critical nodes. This has potential applications in guiding expert elicitation exercises, since such efforts tend to be expensive in time and budget, where results can only be obtained for a limited number of nodes.

References

[1] Eurocontrol, 2004, Review of techniques to support the EATMP safety assessment methodology, Volume I, EEC note no. 01/04.

[2] Iman, Ronald L., 1987, A matrix-based approach to uncertainty and sensitivity analysis for fault trees, *Risk Analysis*, 7(1), pp. 21-33.

[3] Borener, Sherry, S. Trajkov, Poornima Balakrishna, 2012, Design and development of an integrated safety assessment model for NextGen, International Annual Conference of the American Society for Engineering Management.

[4] Shortle, John, Lance Sherry, Arash Yousefi, Richard Xie, 2012, Safety and sensitivity analysis of the advanced airspace concept for NextGen, Proceedings of the Integrated Communication, Navigation, and Surveillance Conference, Herndon, VA, O2-1 – O2-10.

[5] Duijm, N. J., 2009, Safety-barrier diagrams as a safety management tool, *Reliability Engineering and Systems Safety*, 94, pp. 332-341.

[6] Zare, Azin, John Shortle, 2017, Uncertainty analysis for event sequence diagrams in aviation safety. In Transportation Research Board, 96th Annual Meeting, Washington, DC, pp. 1-14.

[7] Noh, Seungwon, John Shortle, 2015, Sensitivity analysis of event sequence diagrams for aircraft accident scenarios, Proceedings of the 34th Digital Avionics Systems Conference, Prague, Czech Republic.

[8] Noh, Seungwon, John Shortle, 2016, Application of common cause failure methodology to aviation safety assessment model, Proceedings of the International Conference on Research in Air Transportation, Philadelphia, PA.

[9] Noh, Seungwon, John Shortle, 2018, Barrier analysis of an aviation safety assessment model. In INCOSE International Symposium, Washington, DC, 28(1), pp. 616-627.

Acknowledgements

The authors thank the support of Aleta Best, Program Manager, System Safety Management Transformation, FAA.

Disclaimer

The opinions expressed in this paper are solely those of the authors.