Bayesian Robustness for Fault Tree Analysis

Chaitanya Joshi
(*with Fabrizio Ruggeri & S.P. Wilson*)

Department of Mathematics & Statistics,

University of Waikato, New Zealand.

13^{th} Nov 2017
BoB 2017, Gold Coast.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Motivation for this work..

HOME ABOUT ACCOMMODATION CALL FOR SUBMISSIONS CONFERENCE REGISTRATION CONTACT
KEY DATES KEYNOTE SPEAKERS LOCATION PROGRAM PROGRAM PRESENTATIONS TRAVEL

BAYES ON THE BEACH 2015 *December 7th – 9th, Surfers Paradise, Queensland*

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Fault Tree Analysis (FTA)

- To evaluate risk in large, safety critical systems.
- To quantify the probability of occurrence of an undesirable event, called the *Top event* (TE).

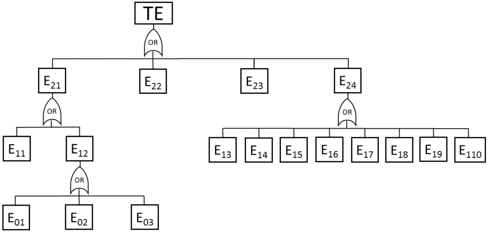

Main assumption: Events are statistically independent.

Bayesian Networks(BN) alternative to FTA

BN can be seen as a natural extension of FTA (FTA can be directly mapped into a BN).

- Can incorporate local dependence between events.
- Enables both forward (prediction) as well as backward (inference) analysis.

However BN approaches require specifying exact prior probabilities for each elementary event and exact conditional probabilities for every dependency!

- Accurate prior probabilities are often not known.
- Computational challenges!

Fully Bayesian implementation of FTA

DePersis (2016)

- Elicit a prior distribution for each elementary event.
- Use simulations to derive the prior distributions for the intermediate events and the TE.
- Find posterior distributions using importance sampling.

Assumes that elementary events are independent.

Prior elicitation for FTA

- Eliciting the (*Beta*) priors using expert opinion is usually not straightforward.
    - Moment matching.
    - Pairwise comparisons (using AHP)
- For very complex systems with little or no data, eliciting even a mean value (of the probability of an event) can be quite challenging for experts.

Elicited priors are likely to be inaccurate!

Prior elicitation is prone to multiple errors!

- Uncertainty due to lack of enough prior information/knowledge.
- Using insufficient or inaccurate information to elicit priors.
- Errors introduced by the methods used for prior elicitation.
- Subjectivity/bias of experts.

Prior mis-specification: snowball effect!

Consider small perturbations to the priors of each of the elementary event.

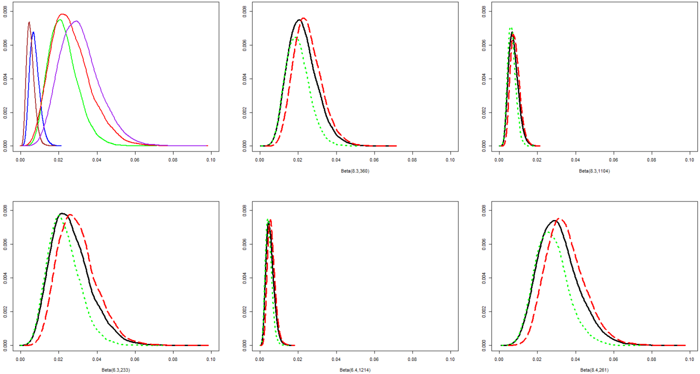

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Prior mis-specification: snowball effect!

The resulting perturbation to the prior of the TE.

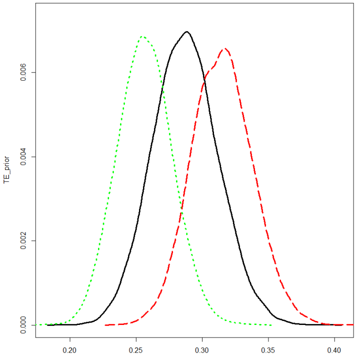

This effect is especially prominent for fault trees containing OR gates.

Posterior influenced by prior

- Data on TE is often sparse and hence posterior is largely determined by prior.
  - TE is often an undesirable event - by definition unlikely to occur (*hopefully*).
  - Safety critical and/or very expensive applications, e.g. spacecraft re-entry!
  - Very little data, if any, available.

Prior mis-specification $\approx$ posterior mis-specification.

Posterior influenced by prior

Posterior(green- dashed) vs prior (black):

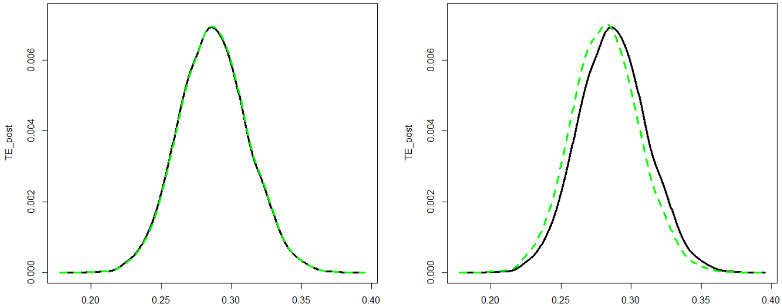

Figure: (*left*) $n = 3$ and 1 failure and (*right*) $n = 10$ and 1 failure.

A distortion function

A *distortion function* h is a non-decreasing continuous function $h : [0,1] \longrightarrow [0,1]$ such that $h(0) = 0$ and $h(1) = 1$. When h is used to transform the distribution function F,

$$F_h(X) = h \circ F(x) = h[F(x)]$$

represents a perturbation of F in order to measure the uncertainty about it. Note that $F_h(X)$ is also a distribution function for a particular random variable denoted by X_h and the distorted density is given by

$$f_h(X) = h'[F(x)] \cdot f(x).$$

Stochastic ordering

For two random variables X and Y, X is said to be *smaller than* Y in the *stochastic order* sense (denoted by $X \leq_{st} Y$) if

$$F_X(t) \geq F_Y(t), \; \forall t \in \mathbb{R}.$$

For absolutely continuous [discrete] random variables X and Y with densities [discrete densities] f_X and f_Y, respectively, X is said to be *smaller than* Y in the *likelihood ratio order* sense (denoted by $X \leq_{lr} Y$) if

$$\frac{f_Y}{f_X} \text{ increases over the union of the supports of } X \text{ and } Y.$$

It is well known that

$$X \leq_{lr} Y \Rightarrow X \leq_{st} Y.$$

Convex and concave distortion functions

- If π is a specific prior belief with distribution function F_π and h is a convex (concave) distortion function in $[0, 1]$, then $\pi \leq_{lr} (\geq_{lr}) \pi_h$.
- If the decision maker is able to represent the changes to a prior belief π by a concave distortion function h_1 and a convex distortion function h_2, then it leads him to two distorted distributions π_{h_1} and π_{h_2} such that $\pi_{h_1} \leq_{lr} \pi \leq_{lr} \pi_{h_2}$.
- This defines the class of priors called the **distorted band of priors $\mathbf{\Gamma_{h_1,h_2,\pi}}$** as

$$\Gamma_{h_1,h_2,\pi} = \{\pi' : \pi_{h_1} \leq_{lr} \pi' \leq_{lr} \pi_{h_2}\}. \tag{1}$$

Arias-Nicolás et al. (2016).

Distortion bands for a prior distribution

Elicited prior (*black*), lower (*green*) and upper (*red*) distortion bands.

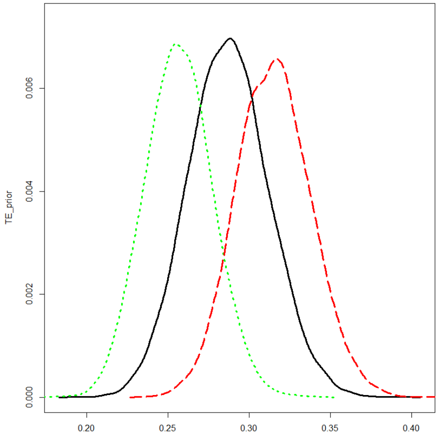

Power functions as distortion functions

A popular choice for distortion functions h_1 and h_2 are power functions given by

$$h_1(x) = 1 - (1-x)^\alpha \text{ and } h_2(x) = x^\alpha,\ \forall \alpha > 1. \tag{2}$$

Note that if we take $\alpha = n \in \mathbb{N}$ in (2), then $F_{\pi_{h_1}}(\theta) = 1 - (1 - F_\pi(\theta))^n$ and $F_{\pi_{h_2}} = (F_\pi(\theta))^n$ which correspond to the distribution functions of the minimum and the maximum, respectively, of an i.i.d. random sample of size n from the baseline prior distribution π.

- Power functions are easily used in applications and also give interesting results.

Power functions - elicitation of α

The Kolmogorov metric measures the maximum absolute difference between the two distribution functions and is defined by

$$K(X,Y) = \sup_{x \in \mathbb{R}} |F_X(x) - F_Y(x)|.$$

Kolmogotov metrics between elicited prior and its distortions are $K(\pi, \pi_{h_1})$ and $K(\pi, \pi_{h_2})$

Interpretation: *How far off could the true prior be from the elicited one in the worst case?*

No more than 20% $\Rightarrow K = 0.2$

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Power functions - elicitation of α

If the distortion functions are defined as in (2) then, the Kolmogorov metric is given by the following expression (Arias-Nicolás et al. (2016)):

$$K(\pi, \pi_{h_1}) = K(\pi, \pi_{h_2}) = \frac{\alpha - 1}{\sqrt[\alpha-1]{\alpha^\alpha}}. \tag{3}$$

Equation (3) can be used for eliciting α.

- Given K, find α using a computer program using (3).
- Alternatively, a rough estimate of α can be obtained assuming $\sqrt[\alpha-1]{\alpha^\alpha} \approx \alpha$ as

$$\alpha \approx \frac{1}{1-K}.$$

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Power functions as distortion functions

For distorted bands obtained using power functions, one can show that:

- α controls the width of the distortion band in a strict monotonic way (so a distorted band obtained using a smaller α is completely contained inside the band obtained using a larger α).
- Stochastic order and likelihood order are equivalent.

Theorem

When the distortion functions are defined as in (2), *(i)*$X \leq_{st} Y \Leftrightarrow X \leq_{lr} Y$, *(ii)* $1 \leq \alpha_1 \leq \alpha_2 \Rightarrow \Gamma_{\alpha_1} \subset \Gamma_{\alpha_2}$ *and (iii)* $\Gamma_\alpha \rightarrow F(x)$ *as* $\alpha \downarrow 1$.

Bayesian robustness for FTA

Given that a prior distribution has been elicited for each of the elementary events.

Bayesian robustness for FTA - an outline

Step I: Build a distorted band of priors for each event.

Step II: Simulate through the FT using algorithms A1 - A4 to find the prior distribution and the distorted band of priors for the intermediate events and the TE.

Step III: Find the posterior distribution for the TE given the prior distribution and the data.

Step IV: Find the lower and the upper distortion bands for the posterior distribution of the TE given the distorted bands for the prior and the data.

Bayesian robustness for FTA

Step I: Build a distorted band of priors for each event i.

- Assume that the prior distribution π_i has been elicited for each of the elementary events.
- Elicit K and therefore α using Equation (3).
- Determine the lower bound $\pi_{h_{1i}}$ using the concave h_{1i} in 2.
- Determine the upper bound $\pi_{h_{2i}}$ using the convex h_{2i} in 2.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Bayesian robustness for FTA

Step II: Simulate through the FT using algorithms A1 - A4 to find the prior distribution and the distorted band of priors for the intermediate events and the TE.

- Algorithm $A1$: to simulate prior distributions for intermediate and top events.
- Algorithm $A2$: to simulate distortion bands for the prior distributions for intermediate and top events.
- Algorithm $A3$: to simulate from $\pi_{h_{1i}}$ for h_{1i} concave.
- Algorithm $A4$: to simulate from $\pi_{h_{2i}}$ for h_{2i} convex.

Bayesian robustness for FTA

Step II: Algorithms $A3$ and $A4$ are rejection sampling based algorithms making use of the fact that h_{1i} (h_{2i}) is concave (convex) and hence has a derivative that is monotonically decreasing (non-decreasing).

Algorithm $A3$: to simulate from $\pi_{h_{1i}}$ for h_{1i} concave

1. Sample $\theta_{ij} \sim \pi_i(\theta_i),\ j = 1, \ldots, N,\ i = 1, 2, 3$ and $u_j \sim U(0, 1)$ independently.
2. For each j, check if $u_j \leq \frac{h_1'[F_i(\theta_{ij})]}{h_1'[0]}$
   - If this holds, accept θ_j as a realisation of $\pi_{h_{1i}}$.
   - If not, reject the value θ_{ij}.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Bayesian robustness for FTA

Step II: Algorithms $A2$ assumes that it is sufficient to sample from $\pi_{h_{1i}}$'s to obtain π_{h_1} and to sample from $\pi_{h_{2i}}$'s to obtain π_{h_2}. It can be proven that this assumption is indeed valid.

Theorem

In order to obtain the distorted lower (upper) bands for the intermediate/top event by sampling from them, it is necessary and sufficient to sample only from the respective lower (upper) bands of the elementary events.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Bayesian robustness for FTA

Steps III and IV: Posterior distribution and the distorted band for the posterior distribution are obtained using the importance sampling algorithm by DePersis (2016).

- Proposal distribution - prior distribution of TE.
- Importance weights using the likelihood.

Example: Spacecraft re-entry

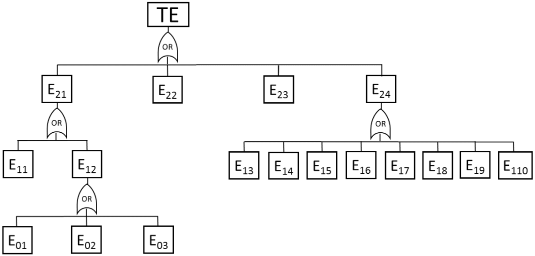

Event	Description	Event	Description
TE	Explosion of the spacecraft	E_{13}	Chemical reactions
E_{21}	Chemical reaction of propellant and air	E_{14}	Over pressure
E_{22}	Burst of pressure vessel	E_{15}	Short circuit
E_{23}	Chemical reaction between hypergolic propellants	E_{16}	Corrosion
E_{24}	Burst of battery cell	E_{17}	Over charge
E_{11}	Sudden release of propellant (E_{22})	E_{18}	Over discharge
E_{12}	Slow release of propellant	E_{19}	Over temperature
E_{01}	Valve leakage	E_{110}	Cell degradation
E_{02}	Tank destruction		
E_{03}	Pipe rupture		

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Example: Spacecraft re-entry

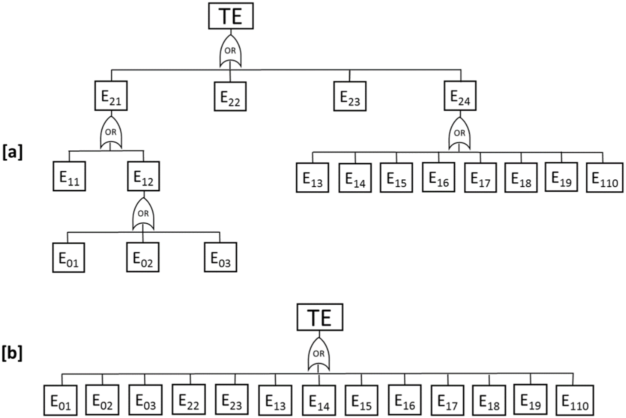

Figure: [a] The fault tree used to model the spacecraft re-entry problem and [b] the simplified fault tree in minimum cut-set

Example: Spacecraft re-entry

Event	Weight	Range	Elicited prior
E_{22}	0.83333	(0.01, 0.05)	Beta(6.3,233) *
E_{23}	0.16667	(0.002,0.01)	Beta(6.4,1214)
E_{01}	0.42857	(0.01, 0.04)	Beta (8.3,360)*
E_{02}	0.1428	(0.0033,0.0133)	Beta(8.3,1104)
E_{03}	0.42857	(0.01,0.04)	Beta(8.3,360)
E_{13}	0.125	(0.014, 0.055)	Beta (8.4,261)*
E_{14}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{15}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{16}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{17}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{18}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{19}	0.125	(0.014, 0.055)	Beta (8.4,261)
E_{110}	0.125	(0.014, 0.055)	Beta (8.4,261)

Table: Elicited priors obtained using the AHP process. * indicates that the prior was elicited using the range provided by the expert

Example: Spacecraft re-entry

Likelihood:

- TE corresponds to whether the spacecraft exploded ($TE = 1$) or not ($TE = 0$) during the re-entry.
- $TE \sim Bernoulli(\theta_{TE})$, where $\theta_{TE} = 1 - \prod_j (1 - \theta_j)$.
- If the data was obtained from n identical spacecraft re-entries then $TE \sim Binomial(n, \theta_{TE})$.
- We assume that only the top event is observed and that none of the elementary events are directly observed.

Distortion bands: We assume that $K = 0.15$ and obtain $\alpha = 1.51$ using Equation 3.

Example: Spacecraft re-entry

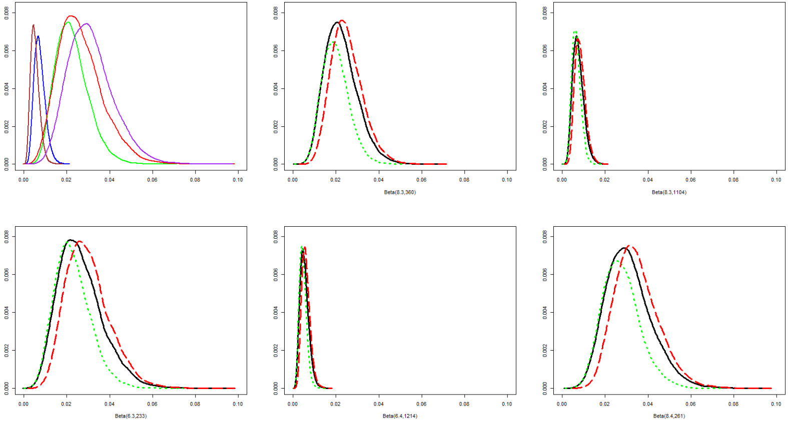

Figure: (*Top left*) the unique prior distributions in Table 1. (*Remaining*) each of the priors and the distorted bands obtained - lower band in *green* - *dotted* and upper band in *red -dashed*.

Example: Spacecraft re-entry

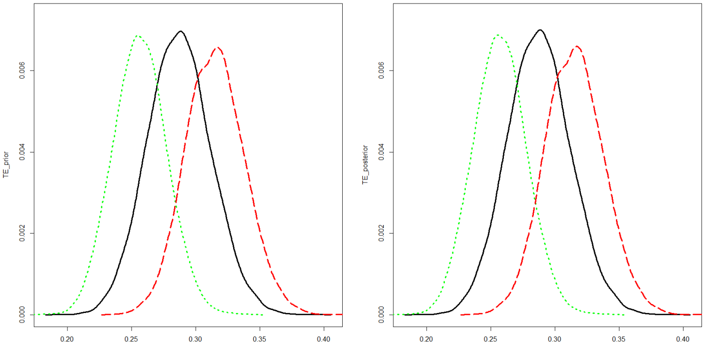

Figure: (*Left*) The prior distribution of θ_{TE} and its distortion bands, (*right*) the posterior distribution of θ_{TE} and its distortion bands: lower band in *green* - *dotted* and upper band in *red -dashed.*

Summary

- This work:
  - Shows how distortion bands obtained using power functions can be used in Bayesian FTA approaches.
  - Provides the sampling algorithms to implement Bayesian FTA.
- Prior robustness study essential!
  - Distortion bands (Arias-Nicolás et al. 2016) have many practical advantages.
- Further/ current work:
  - Prior robustness for ABC methods.