Uncertainty Handling in Integrating Decision Support Systems

Manuele Leonelli

The University of Warwick

CCU2015
January 28, 2015

Decision Making for Nuclear Emergencies

Decision makers (DMs) are nowadays often required to make decisions in *complex* evolving environments.

$\Longrightarrow$ they need to be assisted by a **decision support system (DSS)**.

Decision Making for Nuclear Emergencies

Decision makers (DMs) are nowadays often required to make decisions in *complex* evolving environments.

$\Longrightarrow$ they need to be assisted by a **decision support system (DSS)**.

Traditionally a DSS has often consisted of several *component DSSs*, which guide the estimation and the forecasting of the different relevant quantities.

Each sub-DSS addresses only part of the overall picture.

Decision Making for Nuclear Emergencies

Each component DSS is usually separately built and informed by a separate panel of experts with the relevant knowledge.

Decision Making for Nuclear Emergencies

Each component DSS is usually separately built and informed by a separate panel of experts with the relevant knowledge.

For a potential DM to be properly supported, these disparate judgments need to be *consistently* drawn together to give an overall description of the problem.

Decision Making for Nuclear Emergencies

Each component DSS is usually separately built and informed by a separate panel of experts with the relevant knowledge.

For a potential DM to be properly supported, these disparate judgments need to be *consistently* drawn together to give an overall description of the problem.

A real or virtual system manager, or SupraBayesian (SB), is responsible for the aggregation of these local judgments in order to derive overall expected utility scores.

Decision Making for Nuclear Emergencies

Each component DSS is usually separately built and informed by a separate panel of experts with the relevant knowledge.

For a potential DM to be properly supported, these disparate judgments need to be *consistently* drawn together to give an overall description of the problem.

A real or virtual system manager, or SupraBayesian (SB), is responsible for the aggregation of these local judgments in order to derive overall expected utility scores.

This unifying and integrating framework around which the SB combines component DSSs into a single entity will be called

integrating decision support system (IDSS)

Features of IDSSs

Integration of sub-DSSs into a single Bayesian entity, by stochastically managing the input/output relationships between them:

- It is well know that treating these deterministically, DMs can be directed to choose the wrong course of action (see e.g. Leonelli and Smith 2013).

Features of IDSSs

It is **distributed**, meaning that it is sufficient for each panel to individually deliver their own judgments about the module they oversee:

Features of IDSSs

It is **distributed**, meaning that it is sufficient for each panel to individually deliver their own judgments about the module they oversee:

- The SB can *devolve* the responsibility of each aspect to appropriate experts, making the system more robust to misspecification of beliefs (Cooke 1991);

Features of IDSSs

It is **distributed**, meaning that it is sufficient for each panel to individually deliver their own judgments about the module they oversee:

- The SB can *devolve* the responsibility of each aspect to appropriate experts, making the system more robust to misspecification of beliefs (Cooke 1991);
- *Autonomous* update of beliefs and intervention in the system;

Features of IDSSs

It is **distributed**, meaning that it is sufficient for each panel to individually deliver their own judgments about the module they oversee:

- The SB can *devolve* the responsibility of each aspect to appropriate experts, making the system more robust to misspecification of beliefs (Cooke 1991);
- *Autonomous* update of beliefs and intervention in the system;
- Fast answers to users' queries about the judgments associated with the process and the decisions;

Features of IDSSs

It is **distributed**, meaning that it is sufficient for each panel to individually deliver their own judgments about the module they oversee:

- The SB can *devolve* the responsibility of each aspect to appropriate experts, making the system more robust to misspecification of beliefs (Cooke 1991);
- *Autonomous* update of beliefs and intervention in the system;
- Fast answers to users' queries about the judgments associated with the process and the decisions;
- Potential DMs able to defend their decisions arguing that she has used best expert judgments and the outputs of a *coherent* IDSS to identify her expected utility maximizing policy.

Features of IDSSs

It can be *dynamic*:

- In crisis management it is critical to have the flexibility to allow for algorithmic updates of the relevant probabilities;
- This is because DMs need to make decisions every time new information is gathered throughout the crisis.

The SEU model

Broadly speaking this consists of three main components:

- a *decision space* $\mathcal{D}$ which includes the decisions d available to the DM;
- a *probability density* f over the unknown state $\boldsymbol{y} \in \mathcal{Y}$ of a vector of relevant random variables $\boldsymbol{Y}$;
- a *utility function* $u(d, \boldsymbol{y})$ describing the DM's preferences.

Utility is a real-valued function unique up to positive affine transformations, $u : \mathcal{D} \times \mathcal{Y} \to \mathbb{R}$, such that

$$\forall\ (d_1, \boldsymbol{y}_1), (d_2, \boldsymbol{y}_2) \in \mathcal{D} \times \mathcal{Y}, \quad u(d_1, \boldsymbol{y}_1) \geq u(d_2, \boldsymbol{y}_2) \Leftrightarrow (d_1, \boldsymbol{y}_1) \succeq (d_2, \boldsymbol{y}_2),$$

The SEU model

Broadly speaking this consists of three main components:

- a *decision space* $\mathcal{D}$ which includes the decisions d available to the DM;
- a *probability density* f over the unknown state $\boldsymbol{y} \in \mathcal{Y}$ of a vector of relevant random variables $\boldsymbol{Y}$;
- a *utility function* $u(d, \boldsymbol{y})$ describing the DM's preferences.

Utility is a real-valued function unique up to positive affine transformations, $u : \mathcal{D} \times \mathcal{Y} \to \mathbb{R}$, such that

$$\forall\ (d_1, \boldsymbol{y}_1), (d_2, \boldsymbol{y}_2) \in \mathcal{D} \times \mathcal{Y}, \quad u(d_1, \boldsymbol{y}_1) \geq u(d_2, \boldsymbol{y}_2) \Leftrightarrow (d_1, \boldsymbol{y}_1) \succeq (d_2, \boldsymbol{y}_2),$$

$$\bar{u}(d) = \mathbb{E}(u(d, \boldsymbol{Y})) = \int_{\mathcal{Y}} u(d, \boldsymbol{y}) f(\boldsymbol{y}) \mathrm{d}\boldsymbol{y} = \int_{\mathcal{Y}} u(d, \boldsymbol{y}) \int_{\Theta} f(\boldsymbol{y}|\boldsymbol{\theta}) f(\boldsymbol{\theta}) \mathrm{d}\boldsymbol{\theta} \mathrm{d}\boldsymbol{y}.$$

$$d^* = \arg\max_{d \in \mathcal{D}} \bar{u}(d).$$

Notation and assumptions

- $\boldsymbol{Y}(d) = (Y_1(d), \dots, Y_n(d))$, $d \in \mathcal{D}$;
- Panels of experts $G_1, \dots, G_m$;
- Panel G_i responsible for $\boldsymbol{Y}_i(d) \subset \boldsymbol{Y}(d)$ parametrized by $\boldsymbol{\theta}_i$;
- We call SupraBayesian (SB) the implicit owner of the panels beliefs;

Notation and assumptions

- $\boldsymbol{Y}(d) = (Y_1(d), \ldots, Y_n(d))$, $d \in \mathcal{D}$;
- Panels of experts $G_1, \ldots, G_m$;
- Panel G_i responsible for $\boldsymbol{Y}_i(d) \subset \boldsymbol{Y}(d)$ parametrized by $\boldsymbol{\theta}_i$;
- We call SupraBayesian (SB) the implicit owner of the panels beliefs;

- Variational independence is required i.e.

$$\Theta = \Theta_1 \times \cdots \times \Theta_m$$

Assumptions

- All agree on the class of utility functions $\mathbb{U}$ admitted by the IDSS;
- All agree on the decision rules $d \in \mathcal{D}$ examined by the IDSS;
- All agree the variables $\boldsymbol{Y}$ defining the process, where for each $d \in \mathcal{D}$, each $u \in \mathbb{U}$ is a function of Y, together with a set of qualitative statements about the dependence between various functions of $\boldsymbol{Y}$ and $\boldsymbol{\theta}$ (*structural consensus*);
- $\mathbb{U}$, $\mathcal{D}$ and the structural consensus define the Common Knowledge (CK)- class;
- All agree to take on the beliefs of each panel as their own;

Properties for IDSSs

- **Adequacy**: call an IDSS adequate for a CK-class when the SB can unambiguously calculate her expected utility score, for any decisions $d \in \mathcal{D}$ she might take and any utility function $u \in \mathbb{U}$ she might be given by a user, from the beliefs of panel G_i, $i = 1, \ldots, m$;

- **Soundness**: call an IDSS sound for a CK-class if it is adequate and, by adopting the structural consensus, the SB can coherently admit all the panels' assessments as her own, the SB's underlying belief model being shared with those of the relevant panels.

A simple example

- Two binary random variables $Y_1(d)$, $Y_2(d)$;
- $\theta_1 = P(Y_1 = 1)$, $\theta_{20} = P(Y_2 = 1 | Y_1 = 0)$, $\theta_{21} = P(Y_2 = 1 | Y_1 = 1)$;
- $\boldsymbol{\theta}_2 = (\theta_{21}, \theta_{20})$;

A simple example

- Two binary random variables $Y_1(d)$, $Y_2(d)$;
- $\theta_1 = P(Y_1 = 1)$, $\theta_{20} = P(Y_2 = 1|Y_1 = 0)$, $\theta_{21} = P(Y_2 = 1|Y_1 = 1)$;
- $\boldsymbol{\theta}_2 = (\theta_{21}, \theta_{20})$;
- $\boldsymbol{\mu} = (\mu_{00}, \mu_{10}, \mu_{01}, \mu_{11})$ expectation of $\bar{\boldsymbol{\theta}} = (\bar{\theta}_{00}, \bar{\theta}_{10}, \bar{\theta}_{01}, \bar{\theta}_{11})$ where

  $\bar{\theta}_{00} = (1 - \theta_1)(1 - \theta_{20})$ $\bar{\theta}_{10} = \theta_1(1 - \theta_{21})$ $\bar{\theta}_{01}(1 - \theta_1)\theta_{20}$ $\bar{\theta}_{11} = \theta_1\theta_{21}$

A simple example

- Two binary random variables $Y_1(d)$, $Y_2(d)$;
- $\theta_1 = P(Y_1 = 1)$, $\theta_{20} = P(Y_2 = 1 | Y_1 = 0)$, $\theta_{21} = P(Y_2 = 1 | Y_1 = 1)$;
- $\boldsymbol{\theta}_2 = (\theta_{21}, \theta_{20})$;
- $\boldsymbol{\mu} = (\mu_{00}, \mu_{10}, \mu_{01}, \mu_{11})$ expectation of $\bar{\boldsymbol{\theta}} = (\bar{\theta}_{00}, \bar{\theta}_{10}, \bar{\theta}_{01}, \bar{\theta}_{11})$ where

  $\bar{\theta}_{00} = (1-\theta_1)(1-\theta_{20})$ $\bar{\theta}_{10} = \theta_1(1-\theta_{21})$ $\bar{\theta}_{01}(1-\theta_1)\theta_{20}$ $\bar{\theta}_{11} = \theta_1\theta_{21}$

- suppose $\theta_1 \perp\!\!\!\perp \boldsymbol{\theta}_2$ and let

$$\mu_1 = E(\theta_1) \quad \boldsymbol{\mu}_2 = (\mu_{20}, \mu_{21}) = E(\theta_{20}, \theta_{21})$$

- Then

  $\bar{\mu}_{00} = (1-\mu_1)(1-\mu_{20})$ $\bar{\mu}_{10} = \mu_1(1-\mu_{21})$ $\bar{\mu}_{01}(1-\mu_1)\mu_{20}$ $\bar{\mu}_{11} = \mu_1\mu_{21}$

Panel independence and likelihood separation

- The SB exhibits panel independence if $\perp\!\!\!\perp_{i=1}^{m} \boldsymbol{\theta}_i$;

$$\pi(\boldsymbol{\theta}) = \prod_{i=1}^{m} \pi_i(\boldsymbol{\theta}_i)$$

Panel independence and likelihood separation

- The SB exhibits panel independence if $\perp\!\!\!\perp_{i=1}^{m} \boldsymbol{\theta}_i$;

$$\pi(\boldsymbol{\theta}) = \prod_{i=1}^{m} \pi_i(\boldsymbol{\theta}_i)$$

- A likelihood function is panel separable if

$$l(\boldsymbol{\theta}|\boldsymbol{x}, d) = \prod_{i=1}^{m} l_i(\boldsymbol{\theta}_i|t_i(\boldsymbol{x}, d), d)$$

Panel independence and likelihood separation

- The SB exhibits panel independence if $\perp\!\!\!\perp_{i=1}^{m} \boldsymbol{\theta}_i$;

$$\pi(\boldsymbol{\theta}) = \prod_{i=1}^{m} \pi_i(\boldsymbol{\theta}_i)$$

- A likelihood function is panel separable if

$$l(\boldsymbol{\theta}|\boldsymbol{x}, d) = \prod_{i=1}^{m} l_i(\boldsymbol{\theta}_i|t_i(\boldsymbol{x}, d), d)$$

- Then the posterior

$$\pi^*(\boldsymbol{\theta}) = \prod_{i=1}^{m} \pi_i^*(\boldsymbol{\theta}_i)$$

  where

$$\pi_i^*(\boldsymbol{\theta}_i) \propto l_i(\boldsymbol{\theta}_i|t_i(\boldsymbol{x}, d), d)\pi_i(\boldsymbol{\theta}_i)$$

  and therefore $\perp\!\!\!\perp_{i=1}^{m} \boldsymbol{\theta}_i|\boldsymbol{x}$

Independence models

- Assume $\perp\!\!\!\perp_{i=1}^{m} Y_i | \boldsymbol{\theta}_i, d$;
- let $u(d, \boldsymbol{y}) = \sum_{i=1}^{m} b_i(d) u_i(d, y_i)$;
- then $\bar{u}(d) = \sum_{i=1}^{m} \bar{u}_i(d)$, where

$$\bar{u}_i(d) = \int_{\Theta_i} \bar{u}_i(d, \boldsymbol{\theta}_i) \pi(\boldsymbol{\theta}_i) \mathrm{d}\boldsymbol{\theta}_i$$

  and

$$\bar{u}_i(d, \boldsymbol{\theta}_i) = \int_{\mathcal{Y}_i} b_i u_i(d, \boldsymbol{y}_i) f(\boldsymbol{y}_i | \boldsymbol{\theta}_i, d) \mathrm{d}\boldsymbol{y}_i;$$

- No panel independence needed!!!!

The Multiregression Dynamic Model

The MDM is defined through a DAG G and the following equations: Specifically,

- Observation Equations:

$$\boldsymbol{Y}_t(i) = F_t(i)^T \boldsymbol{\theta}_t(i) + \boldsymbol{v}_t(i), \qquad \boldsymbol{v}_t(i) \sim (0, \Sigma_t(i));$$

- System Equations:

$$\boldsymbol{\theta}_t(i) = G_t(i)\boldsymbol{\theta}_{t-1}(i) + \boldsymbol{w}_t(i), \qquad \boldsymbol{w}_t(i) \sim (0, W_t(i));$$

- Prior Distributions:

$$\boldsymbol{\theta}_0(i) | I_0 \sim (\boldsymbol{m}_0(i), C_0(i))$$

$F_t(i)$ is a function of the parent series;

An example

Consider an MDM defined as

- $Y_2(2) = \theta_2(1,2)Y_2(1) + v_2(2)$;
- $\theta_2(1,2) = \theta_1(1,2) + w_2(2)$;
- $Y_1(2) = \theta_1(1,2)Y_1(1) + v_1(2)$;
- $Y_2(1) = \theta_2(1,1) + v_2(1)$;
- $\theta_2(1,1) = \theta_1(1,1) + w_2(1)$;
- $Y_1(1) = \theta_1(1,1) + v_1(1)$;

- $E(V_t(i)) = \lambda_t(i)$, $E(W_t(i)) = \sigma_t(i)$, $E(\theta_1(1,j)) = a_1(1,j)$ and $V(\theta_1(1,j)) = \tau_1(1,j)$.

An example

Utility factorization:

$$u(y_1(1), y_1(2), y_2(1), y_2(2)) = u(y_1(1)) + u(y_1(2)) + u(y_2(1)) + u(y_2(2)),$$

where $u(y_t(i)) = -y_t(i)^2$.

An example

Overall expected utility $E(u(y_1(1), y_1(2), y_2(1), y_2(2))) =$

$$
\begin{gathered}
-2a_1^2(1,2)\tau_1(1,1) - \tau_1(1,2)(a_1(1,1)^2 + \tau_1(1,1)) \\
-\lambda_1(2) - 2a_1^2(1,2)a_1^2(1,1) - \lambda_1(1)(a_1(1,2)^2 + \tau_1(1,2)) - \lambda_2(1) \\
-(1+\sigma_2(2))(\sigma_2(1) + \lambda_2(1) + \tau_1(1,1) + a_1(1,1)) \\
-(\sigma_2(1) + \lambda_2(1))(\tau_1(1,2) + a_1^2(1,2)) - \tau_1(1,1) - \lambda_1(1) - a_1^2(1,1)
\end{gathered}
$$

An example

Overall expected utility $E(u(y_1(1), y_1(2), y_2(1), y_2(2))) =$

$$
\begin{aligned}
&- 2a_1^2(1,2)\tau_1(1,1) - \tau_1(1,2)(a_1(1,1)^2 + \tau_1(1,1)) \\
&\quad - \lambda_1(2) - 2a_1^2(1,2)a_1^2(1,1) - \lambda_1(1)(a_1(1,2)^2 + \tau_1(1,2)) - \lambda_2(1) \\
&\qquad - (1 + \sigma_2(2))(\sigma_2(1) + \lambda_2(1) + \tau_1(1,1) + a_1(1,1)) \\
&\quad -(\sigma_2(1) + \lambda_2(1))(\tau_1(1,2) + a_1^2(1,2)) - \tau_1(1,1) - \lambda_1(1) - a_1^2(1,1)
\end{aligned}
$$

An example

Overall expected utility $E(u(y_1(1), y_1(2), y_2(1), y_2(2))) =$

$$
\begin{aligned}
&-2a_1^2(1,2)\tau_1(1,1) - \tau_1(1,2)(a_1(1,1)^2 + \tau_1(1,1)) \\
&\quad -\lambda_1(2) - 2a_1^2(1,2)a_1^2(1,1) - \lambda_1(1)(a_1(1,2)^2 + \tau_1(1,2)) - \lambda_2(1) \\
&\qquad -(1 + \sigma_2(2))(\sigma_2(1) + \lambda_2(1) + \tau_1(1,1) + a_1(1,1)) \\
&\quad -(\sigma_2(1) + \lambda_2(1))(\tau_1(1,2) + a_1^2(1,2)) - \tau_1(1,1) - \lambda_1(1) - a_1^2(1,1)
\end{aligned}
$$

Discussion

- IDSSs provide a coherent and efficient framework around which a formal Bayesian decision analysis can be performed;
- a variety of models, both preferential and probabilistic, can be used within the CK-class;
- we developed message-passing algorithm for the distributed computation of expected utilities within IDSSs;
- missing data can break the distributivity of the system and approximated methods need to be developed to address this issue;
- for the purpose of decision making, the algebraic structure of expected utilities can be very helpful;