Aggregating Imprecise Probabilistic Knowledge

Alessio Benavoli and Alessandro Antonucci

IDSIA - "Dalle Molle" Institute for Artificial Intelligence, Lugano (Switzerland)

ISIPTA '09 - Jul. 17th, 2009

Bayesian combination: aggregating sources of precise probabilistic knowledge

- Different sources reporting information about a same variable (X)
- An auxiliary variable for each source (A_1 and A_2)
- Each source returns a conditional mass function ($P_1(X|a_1)$ $P_2(X|a_2)$)
- Each source has a prior ($P_1(A_1)$ and $P_2(A_2)$)
- This defines a joint for each source ($P_1(A_1, X)$ and $P_2(A_2, X)$)
- Conditionals ($P_1(A_1|x)$ and $P_2(A_2|x)$) by Bayes' rule
- Fusion center revises its conditionals from those of the sources
  $P_0(a_1|X) \propto P_1(a_1)P_1(X|a_1)$ $\qquad P_0(a_2|X) \propto P_2(a_2)P_2(X|a_2)$
- Fusion center has its own prior ($P_0(X)$)
- Sources are independent given the variable ($A_1 \perp A_2|X$),
  thus $P_0(A_1, A_2|x) = P_0(A_1|x)P_0(A_2|x)$
- By chain rule, the global joint of the fusion center is therefore
  $P_0(X, A_1, A_2) = P_0(X)P_0(A_1|X)P_0(A_2|X)$
- Bayes' rule $P_0(X|a_1, a_2) \propto P_0(X)P_0(X|a_1)P_0(X|a_2)$
- Flat $P_0(X) \Rightarrow$ Bayesian combination $P_0(X|a_1, a_2) \propto P_1(X|a_1)P_2(X|a_2)$

Graphical outline: from sources to fusion center through model revision

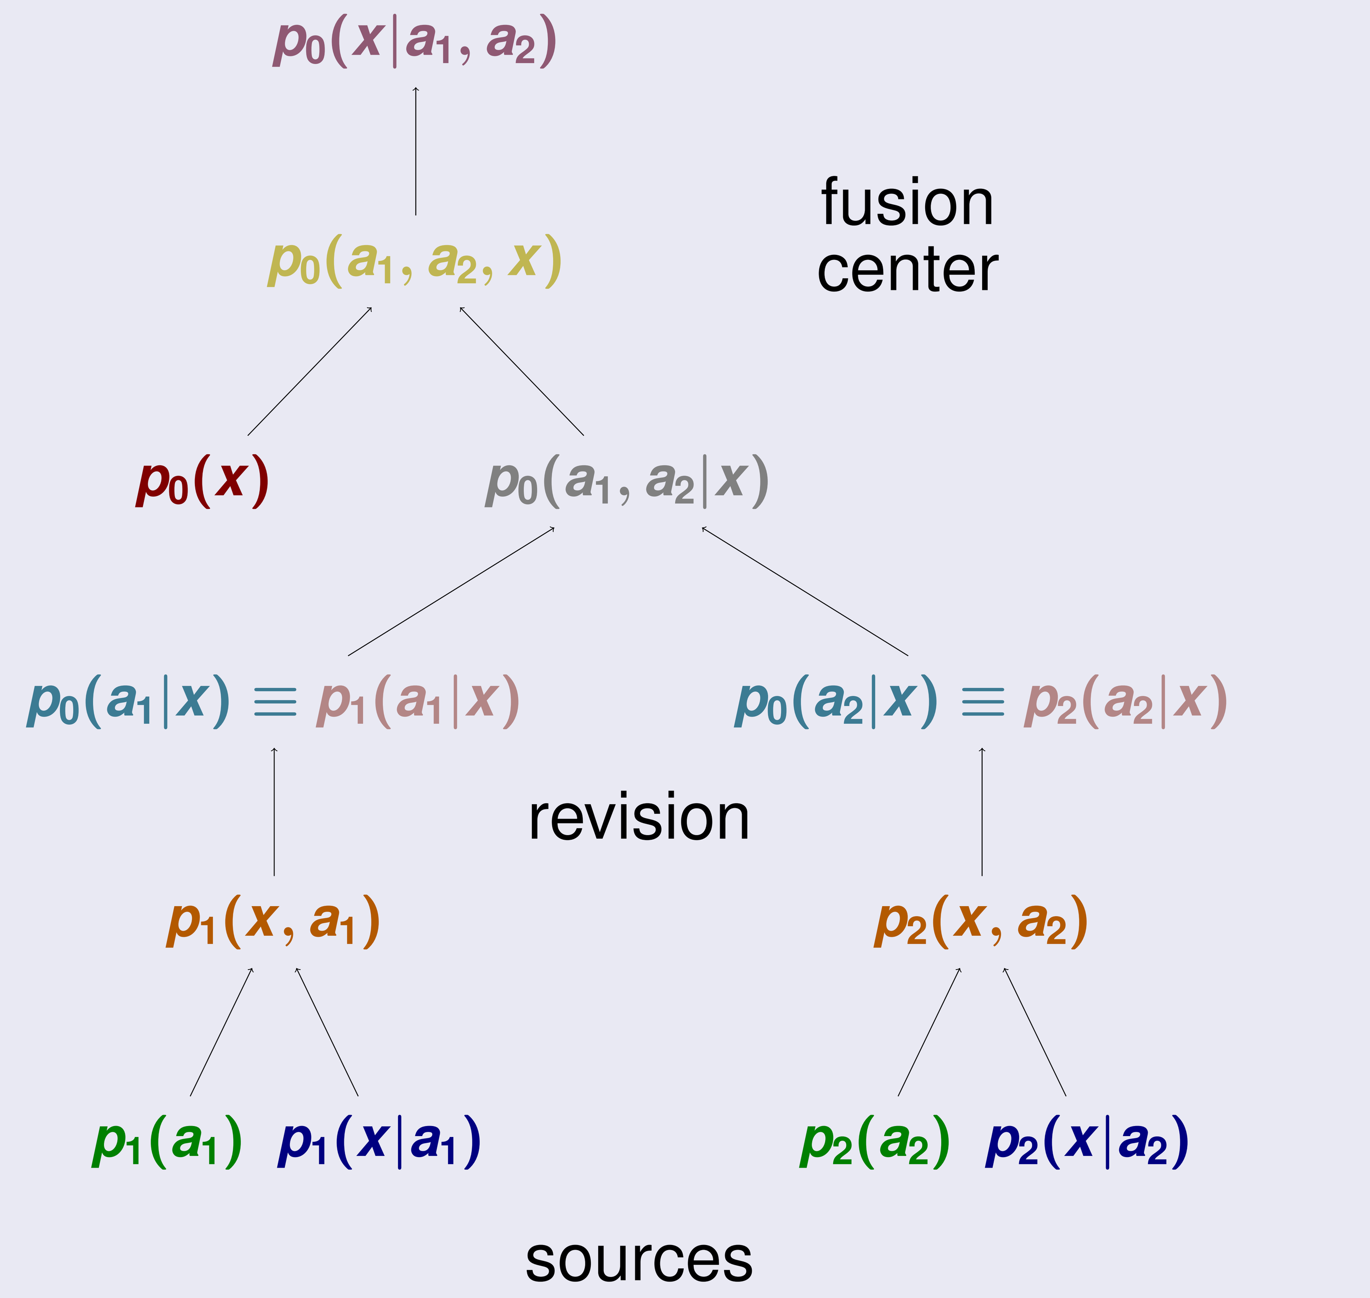

Extension to coherent lower previsions (CLPs): aggregating sources of imprecise probabilistic knowledge

- Each source returns a conditional CLP ($\underline{P}_1^{X|A_1}$ and $\underline{P}_2^{X|A_2}$)
- CLPs modeling also priors for the sources ($\underline{P}_1^{A_1}$ and $\underline{P}_2^{A_2}$)
- Two joints by marginal extension $\underline{P}_j^{X,A_j}(f_j) := \underline{P}_j^{A_j}\left(\underline{P}_j^{X|A_j}(f_j|A_j)\right)$
- Then, conditionals $\underline{P}_1^{A_1|X}$ and $\underline{P}_1^{A_1|X}$ by GBR/regular extension
- Model revision of the corresponding CLPs of the information center
  $\underline{P}_0^{A_1|X} \equiv \underline{P}_1^{A_1|X}$ and $\underline{P}_0^{A_2|X} \equiv \underline{P}_2^{A_2|X}$
- Fusion center has its own prior CLP $\underline{P}_0^X$
- Epistemic irrelevance of the sources given X. Conditional $\underline{P}_0^{A_1,A_2|X}$ by independent natural extension $\underline{P}_0^{A_1,A_2|X}(g|x)$ is

$$\sup_{g_1,g_2} \inf_{a_1,a_2} \left\{g(a_1, a_2) - \left[g_1(a_1, a_2) - \underline{P}_0^{A_1|X}(g_1(\cdot, a_2)|x)\right] - \left[g_2(a_1, a_2) - \underline{P}_0^{A_2|X}(g_2(a_1, \cdot)|x)\right]\right\}.$$

- A joint CLP by marginal extension $\underline{P}_0^{X,A_1,A_2}(g) := \underline{P}_0^X\left(\underline{P}_0^{A_1,A_2|X}(g|X)\right)$
- A separately coherent conditional lower prevision $\underline{P}_0^{X|A_1,A_2}$ by GBR. Assuming $\underline{P}_0^{A_1,A_2}(\tilde{a}_1, \tilde{a}_2) > 0$ ($\tilde{a}_j$ observed internal states of the j-th source), just find μ such that: $\underline{P}_0^{X,A_1,A_2}\left(I_{\{\tilde{a}_1,\tilde{a}_2\}} \cdot [g - \mu]\right) = 0$

Checking coherence

- Sources and fusion center are different *subjects*
- They (asymmetrically) share information by a model revision process
- Coherence required only separately for each subject
- Fusion center: *the separately coherent conditional lower previsions* $\underline{P}_0^{A_j|X}$ *and* $\underline{P}_0^X$ *are jointly coherent* (see proof in the paper)
- Sources: trivial because of marginal extension

A closed formula for linear-vacuous mixtures

- Explicit computation for a special class of CLPs
  $\underline{P}_j^{X|A_j}(f_j|a_j) := \epsilon_j^{a_j} \sum_{x \in \mathcal{X}} p_j(x|a_j)f_j(x, a_j) + (1 - \epsilon_j^{a_j}) \min_{x \in \mathcal{X}} f_j(x, a_j)$
- If the fusion center has a prior vacuous, it will never learn from the sources ($\underline{P}_0^X$ is vacuous $\Rightarrow \underline{P}_0^{X|A_1,A_2}$ vacuous)
- If the sources are vacuous, the fusion center keep its prior as a posterior ($\underline{P}_1^{X|A_1}$ and $\underline{P}_2^{X|A_2}$ vacuous $\Rightarrow \underline{P}_0^{X|A_1,A_2} = \underline{P}_0^X$)
- In general $P(g|a_1, a_2)$ is the (easily computable) solution μ of the following equation:

$$0 = \epsilon_0 \sum_{x\in\mathcal{X}} \left\{\left[\underline{P}_0^{A_1|X}(I_{\{\tilde{a}_1\}}|x) \cdots \underline{P}_0^{A_n|X}(I_{\{\tilde{a}_n\}}|x) I_{\{g(x,\tilde{a}_1,\ldots,\tilde{a}_n)-\mu\geq 0\}} + \overline{P}_0^{A_1|X}(I_{\{\tilde{a}_1\}}|x) \cdots \overline{P}_0^{A_n|X}(I_{\{\tilde{a}_n\}}|x) I_{\{g(x,\tilde{a}_1,\ldots,\tilde{a}_n)-\mu<0\}}\right](g(x, \tilde{a}_1, \ldots, \tilde{a}_n) - \mu)p_0(x)\right\} + (1-\epsilon_0) \min_{x\in\mathcal{X}} \left\{\left[\underline{P}_0^{A_1|X}(I_{\{\tilde{a}_1\}}|x) \cdots \underline{P}_0^{A_n|X}(I_{\{\tilde{a}_n\}}|x) I_{\{g(x,\tilde{a}_1,\ldots,\tilde{a}_n)-\mu\geq 0\}} + \overline{P}_0^{A_1|X}(I_{\{\tilde{a}_1\}}|x) \cdots \overline{P}_0^{A_n|X}(I_{\{\tilde{a}_n\}}|x) I_{\{g(x,\tilde{a}_1,\ldots,\tilde{a}_n)-\mu<0\}}\right](g(x, \tilde{a}_1, \ldots, \tilde{a}_n) - \mu)\right\}$$

Application to Zadeh' paradox

- A patient disease X. Possible diseases are meningitis (x_1), concussion (x_2) and brain tumor (x_3)
- Two sources of information (doctors) $\Rightarrow$ (two Boolean A_1,A_2 s.t., $A_j = a_j$ means doctor j is reliable)
- Doctor A_1 says 99% meningitis, 1% brain tumor, tumor cannot be
- Doctor A_2 says 99% concussion, 1% brain tumor, meningitis cannot be
- After aggregation, either Dempster' rule and Bayesian combination say brain tumor 100%
- In our framework, the joint diagnosis is $\underline{P}_0^{X|A_1,A_2}$
- We have the same result if both the doctors are reliable, i.e., $\underline{P}_0(X|a_1, a_2)$
- But, the conflict says that (at least) one of them should not be reliable
  We can compute $\underline{P}_0(X|\{\neg a_1, a_2\} \cup \{a_1, \neg a_2\} \cup \{\neg a_1, \neg a_2\})$!
  The fusion center conclude that the patient should soffer from either concussion or meningitis