Combination of Bayesian beliefs

Ronnie Johansson

June 29, 2005

Abstract

Derivation of an equation for the combination of Bayesian beliefs.

Proof and discussion

A probability distribution over a set of possible states, Θ, of some mission relevant environment that expresses an individual's or organization's approximative understanding or *belief* of the true state, θ^*, is said to represent *epistemic* uncertainty (as opposed to *aleatory* uncertainty which represents the objective uncertainty that governs a random experiment). In an experiment published elsewhere, we combine the beliefs (i.e, uncertainty grids in that case) of two agents. A cell in observer a_i's grid expresses the observer's belief $P^i_{\mathbf{x}}(\theta)$ that a target sought is located in the corresponding grid position $\mathbf{x} = (x, y)$. Agent a_i's belief is characterized by its history of observations (denoted O_i). Hence it can be written as $P_{\mathbf{x}}(\theta|O_i)$. The result of combining the beliefs of two agents a_j and a_i can now be expressed as $P^i_{\mathbf{x}}(\theta) = P^j_{\mathbf{x}}(\theta) = P_{\mathbf{x}}(\theta|O_i, O_j)$. Note that there is obviously a risk of *double counting*, i.e., using the same observation multiple times (brutely violating any assumption of observation independence) if a pair of agents continually combine their beliefs. However, that discussion and appropriate precautions are beyond the scope of this article.

Neglecting any observation dependencies the resulting combined belief becomes

$$
\begin{aligned}
P_{\mathbf{x}}(\theta|O_i, O_j) &= \{\text{using Bayes' rule}\} = \frac{P_{\mathbf{x}}(O_i, O_j|\theta)P_{\mathbf{x}}(\theta)}{\sum_{\theta' \in \Theta} P_{\mathbf{x}}(O_i, O_j|\theta')P_{\mathbf{x}}(\theta')} = \\
&= \{\text{assuming independence}\} = \frac{P_{\mathbf{x}}(O_i|\theta)P_{\mathbf{x}}(O_j|\theta)P_{\mathbf{x}}(\theta)}{\sum_{\theta' \in \Theta} P_{\mathbf{x}}(O_i|\theta')P_{\mathbf{x}}(O_j|\theta')P_{\mathbf{x}}(\theta')} = \\
&= \{\text{Bayes' rule on } P_{\mathbf{x}}(O|\theta)\} = \frac{\alpha P_{\mathbf{x}}(\theta|O_i)P_{\mathbf{x}}(\theta|O_j)/P_{\mathbf{x}}(\theta)}{\sum_{\theta' \in \Theta} \alpha P_{\mathbf{x}}(\theta'|O_i)P_{\mathbf{x}}(\theta'|O_j)/P_{\mathbf{x}}(\theta')} = \\
&= \{\alpha = P_{\mathbf{x}}(O_i)P_{\mathbf{x}}(O_j) \text{ ind. of } \theta\} = \frac{P_{\mathbf{x}}(\theta|O_i)P_{\mathbf{x}}(\theta|O_j)/P_{\mathbf{x}}(\theta)}{\sum_{\theta' \in \Theta} P_{\mathbf{x}}(\theta'|O_i)P_{\mathbf{x}}(\theta'|O_j)/P_{\mathbf{x}}(\theta')} = \\
&= \{\text{assuming } P_{\mathbf{x}}(\theta) \text{ uniform}\} = \frac{P_{\mathbf{x}}(\theta|O_i)P_{\mathbf{x}}(\theta|O_j)}{\sum_{\theta' \in \Theta} P_{\mathbf{x}}(\theta'|O_i)P_{\mathbf{x}}(\theta'|O_j)} .
\end{aligned}
\tag{1}
$$

Note that the end result in Equation 1 is an expression that contains only the known beliefs of agents a_i and a_i. It is also interesting to note that the result is consistent with Dempster's rule (from evidential theory) when the belief functions only have singletons as their focal elements.

It is important to remember that this result is based on two assumptions: i) a common and uniform initial prior $P_{\mathbf{x}}(O)$ for all agents, and ii) observation independence. If the assumption of a uniform prior is dropped, then the initial prior must be kept and used in all beliefs combinations. If also the assumption of a common prior is dropped, the agents might get different posteriors, i.e., $P^i_{\mathbf{x}}(\theta|O_i, O_j) \neq P^j_{\mathbf{x}}(\theta|O_i, O_j)$. Assumption ii) is highly unrealistic, but frequently accepted for practical reasons.

Finally, note that the calculation of Equation 1 requires that $P_{\mathbf{x}}(\theta) > 0$. If $P_{\mathbf{x}}(\theta)$ happens to be zero for some θ, it is reasonable (and consistent with Bayesian updating) to conclude that also the posterior is zero, i.e., $P_{\mathbf{x}}(\theta|O_i, O_j) = 0$.