Composing Bayesians: A Short Review

Robin Hanson

May 14, 1991

Why Compose?

One of our best techniques for dealing with software (or any kind of) complexity is modularity. We get a lot of mileage out of being able to combine pre-existing software modules into new computer programs, especially when we can predict the properties of those new modules from the properties of their parts. For example, if the parts are fast, the total can be fast. If the parts are "correct" and are combined correctly the total is correct. Since being Bayesian is supposed to be a preferred property of agents, including computer-based ones, the question naturally arises: if we combine programs that are "Bayesian", can we get a new program that is also Bayesian?

It's a Basic Question

This issue has been widely discussed in the context of statistics [Bordley, 1986; French, 1985; Genest and Zidek, 1986] and philosophy, where one wonders whether a group of Bayesian (or "rational") agents can itself be Bayesian, and in political economy, where one asks how citizens can combine their votes into a collective decision. It is also relevant in cognitive science where one wonders what properties hold for the modules that make up our exceedingly complex mind. In general we suspect that if being Bayesian is good for a unit at one level of abstraction, then it may well be good for units at other levels of abstraction also. The following is a review of some proposed methods for composing Bayesians.

Let's Make a Deal

What if we just let agents with conflicting goals and beliefs make whatever deals they can with each other? It is often assumed that this only leads to chaos, and thus the agent's interactions must be structured somehow (more about this later).

To Each Her Own (Specialty)

We can compute on a Bayes Net by having many specialist agents, each responsible for enforcing the probability constraints associated with a particular node in the network [Pearl, 1988]. But what if the agents have overlapping specialties and disagree? What if they can't even agree on the network?

Herr Dictator

We can make one special "supervisor" agent which considers the recommendations of the other agents by just treating them as evidence. If we take the beliefs and goals of this Bayesian leader to be those of the group, then of course the group is Bayesian. But the leader must now have a (possibly very complex) model of the advisors and their correlations, describing how likely they are to give different advise in different states of the world. And if the non-leader agents have a say, they may not agree to this arrangement.

Vote On It

What about democracy? Regarding the various possible actions the group as a whole could take, let's say that each agent can prefer action A to B, B to A, or be indifferent between them. We could ask each agent what their preferences are regarding the possible group actions, and compute the actual group action as a function of these responses. There are several features we might like such a "social choice function" to have:

- *No Dictator* – There's no individual who always gets her way, even when every one else disagrees.
- *Transitive* – The group preferences are transitive in both preference and indifference
- *Respects Unanimity* – If at least one person prefers A to B, and no one prefers B, the group prefers A.
- *Universal* – Works for any set of (transitive) individual preferences.
- *Local* – What the group thinks about A vs. B only depends on what the members think about these two options, ignoring all other options.

Arrow's impossibility theorem [Mueller, 1979] say's that you can't have a voting rule with all these features. For example, the usual majority voting rules violate transitivity of preference. (If we don't require transitivity of indifference, a social choice function is possible.)

Can Bayesian's Vote?

A Bayesian voting rule for Bayesian citizens must be transitive, but needn't be universal or local. Unfortunately, this doesn't save us. In general, a voting rule which combines Bayesians, has no dictator, and respects unanimity can't result in a Bayesian group [Seidenfeld, 1990]. However there are exceptions. If the agents all agree on probabilities, and if their utilities can be made "comparable", then a linear combination of the utilities is Bayesian.

Linear Opinion Pools

Similarly, if we assume the agents have the same utilities, we can construct a group belief $g(A)$ in any proposition A as a weighted combination of the individual beliefs $p_i(A)$

$$g(A) = \sum_i w_i p_i(A)$$

with $\sum_i w_i = 1$. This rule has the advantage of being uniform (i.e., for all A $g(A)$ is the same function of only the $p_i(A)$), and is only certain (i.e., 0 or 1) when all the agents are certain. As given, the linear pool is only statically Bayesian. If the group belief is to be updated in a Bayesian way on new evidence E, we must change the weights $w_i \to w'_i$ depending on how well each agent predicted that evidence.

$$\frac{w'_i}{w_i} = \frac{p_i(E)}{g(E)}$$

This approach is equivalent to a dictator who uses a linear mixture model of her advisors, basically assuming that one advisor is "right" and trying to infer which one that is.

Logarithmic Opinion Pools

A similar approach is pick a particular partition of events x, define group beliefs on this partition to be

$$\log(g(x)) \propto \sum_i w_i \log(p_i(x))$$

and define other beliefs in terms of these. This approach is not uniform, and results in the group assigning zero probability if any member does. But the group does update its beliefs in a Bayesian way without needing to change the member's weights.

Who's More Expert?

To pick the weights w_i equitably, we could use a self-consistency criterion

$$w_i = \sum_j W_{ij} w_j,$$

where W_{ij} is the weight agent j would assign to agent i. Each agent's consensus weight is then an average of what other's think of them, weighted by the consensus weights of those other agents.

Opinion pool weights cannot vary from issue to issue, but agent's can represent their relative expertize in different subjects by using the group belief as evidence in choosing their own beliefs, and relying more heavily on the group in subjects away from their specialty. Doing this consistently can be tricky though.

Pools are Fragile

To be Bayesian, opinion pools require that all participating agents always be Bayesian, that they all update their beliefs on the same evidence, and that they are honest in declaring what they believe. These assumptions are not always realistic, especially when the agents are competing or distrust on another.

Inferring Evidence

The group might be better off informationally if all agents shared all evidence with each other. And in principle, if each agent knows what kind of information other agents have, they can infer this information from those agent's betting offers [Sebenius and Geanakoplos, 1983]. But with all this redundancy of inference, it's not clear what the point of having multiple agents would be, computationally speaking. Instead, all of the approaches in this review rely to some degree on agents making more limited inferences from what other agents say they believe, though for some approaches this is only an optimization issue. Thus we should not expect all agents to be considering all evidence.

Buying Honesty

Non-local voting rules can create incentives for "strategic voting", such as lying about preferences and trading votes. How can we promote honesty?

If a participant in a linear pool has a utility linear in $\log(\frac{w_i}{1-w_i})$, then it is in that participants self-interest to be honest [Bayarri and DeGroot, 1988]. If, however, we just know that an agent has a utility monotonic in something like money, then we might try to buy honesty by making bets or offering a small monetary reward linear in a "scoring rule" like $\log(p_i(E_j))$ where E_j is the event that actually occurs. This induces honesty *if* the agent has no other stakes in the question. It turns out, however, that the agents are better off if they make a deal whereby they all claim to believe the same thing and divide up the winnings depending on the situation.

Bettors Appear to Agree

More generally, if the agents are free to bet with one another, then they should create a betting market and each bet as much as they want at the market odds. At equilibrium, no one will want to bet anymore, and so they will each have the same marginal value for more reward conditional upon any of the bet-upon events [Kadane and Winkler, 1988]. Thus to an external observer who can offer rewards or observe external actions, but who cannot tell how much each agent have bet, the agents act as if they agree. Finite transaction costs and time delays add small corrections to this, but the basic result remains.

This result can be understood from the basic expected utility equation

$$EU_i(a) = \sum_s U(\$ \text{ for } i \text{ if } s) p_i(s|a)$$

where the s are the possible states of the world and a is a possible action. While the $p_i(s|a)$ can vary with agent i, the "$ for i if s" can vary to compensate, so that the different agents prefer the same actions.

The Market Consensus

If we give up trying to keep the agents honest, we can just accept the equilibrium market odds as a spontaneously created group belief, somewhat similar to a blackboard as a computational mechanism. This consensus should be Bayesian since any violation of the Bayesian axioms creates an opportunity for any agent to make a riskless profit through an appropriate combination of bets. This result does not depend on all or even any of the participant agents being Bayesian, just on being able to detect and exploit profit opportunities. Evidence seen by only a fraction of the agents can be reflected in a Bayesian update of the consensus, since those agents can confidently make very large bets. Specialists can just focus on specialty areas, or they can overlap and disagree, as they choose. Of course there are overhead costs to creating markets, possible strategic games with small numbers of traders, and one needs a way to settle each bet. But overall, just letting agents make deals without constraint is not as bad as it seems.

References

[Bayarri and DeGroot, 1988] M. Bayarri and M. DeGroot. Gaining weight: A Bayesian approach. In Bernardo, DeGroot, Lindley, and Smith, editors, *Bayesian Statistics 3*, pages 25–44. Oxford University Press, 1988.

[Bordley, 1986] Robert Bordley. Review essay: Bayesian group decision theory. In B. Grofman and G. Owen, editors, *Information Pooling and Group Decision Making*, pages 49–68. JAI Press Inc., London, 1986.

[French, 1985] Simon French. Group consensus probability distributions: A critical survey. In Bernardo, DeGroot, Lindley, and Smith, editors, *Bayesian Statistics 2*, pages 183–201. Elsevier Science Publishers, 1985.

[Genest and Zidek, 1986] C. Genest and J. Zidek. Combining probability distributions: A critique and annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.

[Kadane and Winkler, 1988] J. Kadane and R. Winkler. Separating probability elicitation from utilities. *Journal of the American Statistical Association*, 83(402):357–363, June 1988.

[Mueller, 1979] Dennis Mueller. *Public Choice*. Cambridge University Press, 1979.

[Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, 1988.

[Sebenius and Geanakoplos, 1983] J. Sebenius and J. Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.

[Seidenfeld, 1990] Teddy Seidenfeld. Two perspectives on consensus for (Bayesian) inference and decisions. In H.E. Kyburg, editor, *Knowledge Representation and Defeasible Reasoning*, pages 267–286. Kluwer Academic, 1990.