

Leveraging Consensus and Divergence in Bayesian Belief Aggregation

Kshanti Auster Greene

Advisor: George F. Luger

University of New Mexico, Department of Computer Science

1 Research Summary

Many fields have a need to build representative or predictive models from a number of unique individuals who each can contribute their experience and beliefs to the whole. For instance, intelligence agencies may wish to build a model from a number of experts to analyze potential terrorist attacks. In addition, a sociological survey may want a model representing the beliefs of cultural or political groups. However, challenges remain that have limited the success of merging opinions to form consensus models. Our research in progress presents a new approach to combine, or *aggregate* the beliefs of many individuals using graphical models. Existing Bayesian belief aggregation methods utilize an *opinion pool* function to find a single consensus on a given probability distribution [Matzkevich and Abramson, 1992; Pennock and Wellman, 1997; Napp and Jouini, 2006; Maynard-Reid and Chajewska, 2001]. These opinion pool functions have many theoretical problems including breaking several assumptions for Bayesian reasoning [Pennock and Wellman, 1999]. More practically, existing opinion pool functions do not represent reality well, especially in cases of diverse opinions.

Divergence is a natural result of combining opinions from individuals with different beliefs, backgrounds and experiences. Divergence could arise for many reasons including, but not limited to, inaccuracies in the sources. Since there are many valid reasons why disagreement occurs, our aggregation approach aims to retain this divergence. Graphical models such as Bayesian networks require a probability distribution for each variable (node) and relationship (formed by edges) in the model [Pearl, 1988]. Existing opinion pool functions typically form a single consensus value for each probability distribution that will *average out* any divergence. Instead of forcing a single consensus, we leverage both consensus and divergence to find clusters of belief formed by groups of individuals that agree on a given probability distribution. Multiple clusters will form in situations in which there is significant variance in opinion on a probability distribution.

While our belief clusters or *groups of consensus* provide a more realistic representation of aggregate belief, they also present an additional challenge. Before one can extract a prediction from a graphical model, inference must be performed on the model. Inference is the process of propagating evi-

dence and probability estimates throughout a model to form a probability estimate for a given hypothesis [Pearl, 1988]. Instead of propagating a single aggregated consensus value from the opinion pool function throughout the network, we now must propagate multiple values for each probability distribution formed by the belief clusters. The probability of a given node is a function of the conditional probabilities from the node's parents, and the posterior probability of each of its parents. Given k_i clusters and n probability distributions, we need to propagate an exponential number of values throughout the network because each k_i values at each node will need to be combined with the k_j values from each of a node's parents.

To reduce this combinatorial explosion of values during propagation we will leverage an observation that individuals that agree on one issue often agree on a number of issues. Our approach is to find the groups that often agree and propagate the consensus probability distributions throughout the network using these *threads* of agreement.

To discover groups of consensus we utilize social networks to model the agreement between individuals across multiple probability distributions. We form *agreement networks* in which nodes represent individuals and an edge between two nodes indicates that the individuals agreed on one or more beliefs. The weight of the edge indicates the number of beliefs on which the individuals agree. We then use a graph cut algorithm [Shi and Malik, 2000] to discover the groups of agreement in the network. The graph cut algorithm selects edges to cut between clusters of nodes by minimizing the weight of the edges which are cut. The benefit of the graph cut algorithm over other clustering approaches is that it does not require an initial k value for the number of clusters, and it works on small samples as well as large.

Given an arbitrary Bayesian network, if we were to apply the graph cut algorithm using all beliefs in the network, we may find that we do not find a clean *separation* between clusters. Separation is a function of the distance between the mean of the clusters and the variance within the clusters. As the number of beliefs increases, the likelihood that individuals remain in the same groups of consensus decreases, resulting in reduced separation. Therefore, we will first attempt to find a partitioning of the network such that the graph cut algorithm is applied across beliefs that result in higher separation. We discover this partitioning of the network by iter-

actively computing the graph cut on sets of adjacent beliefs, then computing the distance between the means of the clusters and the variance of the clusters.

We would like to balance the separation between clusters with the number of subnetworks in our network partitioning. A smaller number of subnetworks implies a reduced number of times values must propagate to other subnetworks, therefore reducing combinatorial explosion. However, as we increase the size of our subnetworks, we also reduce the separation between clusters. To find a good balance, we will normalize the separation (Δ) by a to-be-determined variable- for instance, the number of sub-graphs (S). Finally, we then run inference in each subnetwork using a single set of clusters for the whole subnetwork, reducing the explosion of values during inference compared to assigning each belief a unique set of clusters.

Our primary goal is to build a consensus model that best represents the divergent opinions while constraining the complexity that results from maintaining diversity. Determination of success and completion of the system is based on a number of metrics. First, the aggregation method should enable the merging of beliefs such that the estimates from members of a consensus cluster are within a specified (maximal) variance. Second, we would like to control the combinatorial explosion that could be introduced by propagating multiple means for each probability distribution throughout the network. Our final goal is to claim that our approach more realistically represents the opinions of divergent groups than existing approaches. This will be measured by comparing the distribution of the result of querying each probability estimate using the aggregate beliefs with the distribution of the query result using individual belief estimates (in which each individual has its own belief network).

Our research will advance the state of Bayesian reasoning because it will enable researchers to build more realistic models from a diverse set of experts, laypeople and software agents. Our research also has long-reaching implications in the areas of collective intelligence, data fusion and collaborative filtering. For instance, agreement across multiple beliefs may indicate a similar background between the individuals. Beliefs may also be an indicator of expertise in certain areas, which can then be used to quantify the quality of a new belief estimate. Our approach will also provide the basis to study the effect of diversity and expertise on the quality of knowledge and predictive models.

2 Research Plan and Progress to Date

I successfully defended my thesis proposal in September, 2008 and was recently awarded candidacy. Since then I have concentrated the focus of my approach by leveraging existing techniques for Bayesian reasoning that enable me to concentrate on immediate issues in belief aggregation. For instance, the noisy OR-gate combining rule [Pearl, 1988] will enable me to utilize independence assumptions and reduce complexity. I also incorporated the normalized graph cut [Shi and Malik, 2000] used primarily in image analysis and vision to find the groups of consensus in the agreement network. Using a dataset of political beliefs I elicited using Amazon's Mechan-

ical Turk (mturk.com), I successfully demonstrated a situation in which a graph cut algorithm discovered considerable agreement across beliefs. I presented my approach and preliminary results at the 2009 AAAI Spring Symposium in the Technosocial Predictive Analytics workshop. My remaining tasks and plan for completion are as follows (all dates are 2009 except where indicated):

January: Develop graph cut algorithm to discover groups of consensus in an agreement network.

February: Implement the noisy OR-gate algorithm to compute the probability of a variable given its parents.

March-May: Define the goals of the network partitioning algorithm and implement.

June: Finalize set of datasets for testing.

July-August: Research and develop the algorithm to propagate posterior probabilities between subnetworks.

September - December: Compile results in thesis.

January 2010: Defend thesis.

3 Acknowledgements

We would like to thank Joseph M. Kniss at UNM and Carl Stern of Management Sciences, Inc. and for their invaluable contributions to this research.

References

- [Matzkevich and Abramson, 1992] Matzkevich, I., and Abramson, B. 1992. The topological fusion of bayes nets. In *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*.
- [Maynard-Reid and Chajewska, 2001] Maynard-Reid, P. I., and Chajewska, U. 2001. Aggregating learned probabilistic beliefs. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*.
- [Napp and Jouini, 2006] Napp, C., and Jouini, E. 2006. Aggregation of heterogeneous beliefs. *Journal of Mathematical Economics* 42(6):752–770.
- [Pearl, 1988] Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman.
- [Pennock and Wellman, 1997] Pennock, D. M., and Wellman, M. P. 1997. Representing aggregate belief through the competitive equilibrium of a securities market. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*.
- [Pennock and Wellman, 1999] Pennock, D. M., and Wellman, M. P. 1999. Graphical representations of consensus belief. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*.
- [Shi and Malik, 2000] Shi, J., and Malik, J. 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8).