Satisficing the masses: Applying game theory to large-scale, democratic decision problems

Kshanti A. Greene, Joseph M. Kniss, George F. Luger
Department of Computer Science
University of New Mexico
Albuquerque, NM 87131
Email: [kshanti, jmk, luger]@cs.unm.edu

Carl R. Stern
Management Sciences, Inc.
6022 Constitution Ave. NE
Albuquerque, NM 87110
Email: carl_stern@mgtsciences.com

I. Introduction

We present ongoing research on large-scale decision models in which there are many invested individuals. We apply our unique Bayesian belief aggregation approach to decision problems, taking into consideration the beliefs and utilities of each individual. Instead of *averaging* all beliefs to form a single consensus, our aggregation approach allows divergence in beliefs to emerge. In decision models this divergence has implications for game theory– potentially turning a cooperative situation into a competitive one. By applying our approach to the topical issue of stem cell research using input from many diverse individuals, we analyze the behavior of a decision model including the groups of agreement that emerge. We discuss the issues involved in finding Nash equilibrium and minimax solutions. We analyze a range of outcomes between satisficing in an attempt to "please everybody," to showing the effect of optimizing the outcome for a small group of individuals. Our approach has the long-reaching potential to help define policy and analyze the effect of policy change on individuals.

II. Background

We base our research on a framework that is well-studied in Artificial Intelligence. *Bayesian networks*, also known as *belief networks*, are a form of graphical model that integrate the concepts of graph theory and probabilistic reasoning [1]. These networks define dependencies between variables that can represent causality, implication or correlation. In a typical Bayesian network, random variables are represented by nodes and conditional relationships are represented by directed edges between the nodes. A variable is *conditioned on* all of its parents, described by the expression $P(X|Pa_x)$ where Pa_x is the set of parents of X.

Bayesian networks can be extended to deal with decision problems using *influence diagrams* [2]. In addition to nodes representing random variables (or *chance* nodes), influence diagrams contain *decision* nodes, representing a decision to be made; and *utility* nodes, representing the value or risk associated with a possible outcome. Influence diagrams efficiently represent the uncertainty involved in real-world decision problems.

Bayesian belief aggregation is the process of combining probability estimates from multiple human or software agents. Belief aggregation typically uses an *opinion pool* function to form a single aggregate distribution from multiple beliefs. However, researchers have shown that it is not possible to maintain consistent structures using an opinion pool function on conditional probabilities [3].

Belief aggregation raises a more philosophical issue that has thus far not been discussed in the literature. The logic behind averaging to find one *consensus* based on many possibly divergent opinions is flawed. Consider the following situation; Joe believes that Democrats winning the election is very unlikely (10%). Susan believes that Democrats winning the election is almost certain (90%). The result of averaging these opinions implies that people believe the election is a tossup, while the individual opinions clearly are quite polarized. A second situation has the opinions; 55% and 45%. The average of these also calls the election a tossup, but the opinions more closely reflect this conclusion. To maintain a realistic representation of belief the resulting consensus model should distinguish between these two situations.

To our knowledge, belief aggregation has not been applied to influence diagrams or decision problems. In this paper, we apply our unique approach to influence diagrams, introducing some interesting game-theoretic implications for analyzing decision problems involving many individuals with diverse beliefs and motivations.

III. Our belief aggregation approach

Our research presents a new approach to combine the beliefs of many individuals using graphical models.

Existing Bayesian belief aggregation methods break theoretical assumptions for Bayesian reasoning and do not generate a realistic representation of diverse opinions. Divergence is a natural result of combining opinions from individuals with different beliefs, backgrounds and experiences. Our approach leverages agreement and disagreement between individuals to reduce the error that occurs during aggregation, as well as to form a more representative consensus model. Instead of computing a single consensus value (or average) to represent the beliefs of many potentially divergent opinions, we cluster similar probability estimates to form *consensus belief clusters* and apply an opinion pool function to each of the clusters. The result is a set of k distributions representing the clusters.

Bayesian inference is the process of propagating probability distributions across network nodes to compute the overall joint probability distribution of the variables in a network. When using our belief clusters instead of one consensus value for each probability distribution, we have a set of k distributions that need to be propagated, resulting in a combinatorial explosion during inference. To reduce this explosion we leverage the agreement that occurs across sets of beliefs. More specifically, we cluster on beliefs between which there is a high covariance.

IV. Application to decision models

We extend the concept developed in [4] in which a multi-agent influence diagram (MAID) forms a model of a competitive situation between two agents. In their MAIDs, the authors define a *decision rule* to be the choice for decision D that is selected given an instantiation of the parents $Pa(D)$, where $Pa(D)$ is the set of variable nodes that effect the decision. A *strategy profile* is an assignment of decision rules to each decision in the MAID. An optimal strategy profile for an agent is one in which the agent's expected utility is maximized for all decisions. A strategy profile is a Nash equilibrium if for all agents in a MAID, their strategies are optimal assuming no other agent changes its strategy.

The MAID in [4] assumes that an agent represents a single entity. In our case, we consider the aggregated beliefs and utilities of potentially many individuals. If a situation existed in which individuals consistently agreed with the same individuals across all beliefs and utilities in the influence diagram, then we could have a single agent represent each belief cluster and find the optimal decision strategies using the MAID approach. However, as in our belief networks, individuals will not always agree with the same individuals, complicating the optimization strategies. In this paper we will discuss the application of our aggregation approach to decision networks and the implications it has for finding optimal strategies for the individuals involved.

V. Game theoretic analysis

Our approach has broad implications for decision-making, particularly in democratic situations in which many diverse individuals have a stake in the outcome. For example, a policy-maker could generate a model to determine a Nash equilibrium decision strategy based on a representative distribution of opinions. A policy-maker could also attempt to improve the worst case scenario by finding a minimax solution. We show that in enabling the divergence in opinions to emerge, we can distinguish a competitive situation from a cooperative situation in which all beliefs form a single consensus.

VI. A decision model for stem cell research

In this paper we apply our approach to a topical decision that is currently being addressed by policy-makers. The issue of stem cell research (Figure 1) is particularly appropriate because of its polarizing effect on individuals with diverse backgrounds and motivations. We will survey individuals using web tools such as mechanical turk (mturk.com) and facebook (facebook.com) and analyze the resulting aggregate models. We are particularly interested in individuals with conflicting motivations, for instance a person of strong Christian faith who also has a child that might benefit from embryonic stem cell research. These individuals represent the outliers that defy attempts to aggregate on the basis of all beliefs. Our approach that aggregates across subsets of more correlated beliefs allows these outliers to emerge.

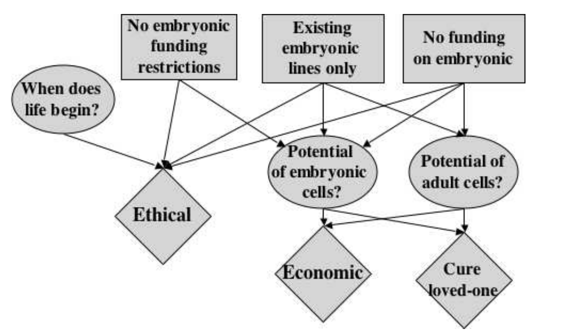

Fig. 1. A partial influence diagram for the stem cell research issue.

References

[1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, 1988.

[2] R. Howard and J. Matheson, Eds., *Readings on the Principles and Applications of Decision Analysis*. Strategic Decisions Group, 1984, vol. 2, pp. 6–16.

[3] D. M. Pennock and M. P. Wellman, “Graphical representations of consensus belief,” in *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 1999.

[4] D. Koller and B. Milch, “Multi-agent influence diagrams for representing and solving games,” in *IJCAI*, 2001, pp. 1027–1036.