

# 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. ABSTRACT

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 and utilities to emerge. In decision models this divergence has implications for game theory—enabling the competitive aspects in an apparent cooperative situation to emerge. Current approaches to belief aggregation assume cooperative situations by forming one consensus from diverse beliefs. However, many decision problems have individuals and groups with opposing goals, therefore this forced consensus does not accurately represent the decision problem. 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 show how to find the Pareto optimal solutions, which represent the decisions in which no group can do better without another group doing worse. We analyze a range of solutions, from attempting to “please everybody,” with the solution that minimizes all emerging group’s losses, to optimizing the outcome for a subset of individuals. Our approach has the long-reaching potential to help define policy and analyze the effect of policy change on individuals.

## II. INTRODUCTION

*As citizens of a vibrant and varied democracy, how do we engage in vigorous debate? How does each of us remain firm in our principles, and fight for what we consider right, without*

<sup>1</sup>A version of this article is being published by IEEE Computer Society

*demonizing those with just as strongly held convictions on the other side?* – United States President Barack Obama

In the previous quote, United States President Obama describes his wish to enable collaborative decision-making while maintaining the diversity of opinion of all individuals with a stake in the outcome of the decision. While computer scientists do not typically delve into politics, we see the potential of large-scale, social decision models to help achieve the President’s goals and the goals of those who wish to work together to solve some of the world’s problems, both small and large.

Our long-term vision is to develop a new social decision-making paradigm that enables people to visualize how their actions affect their own circumstances and their environment, taking into consideration the simultaneous actions and goals of other community members. Our approach will help collective members choose the best way forward based on models formed from their beliefs, motivations and preferences, as well as the actions taken by likeminded individuals. It will also help policy-makers analyze how their far-reaching decisions affect the community they represent, taking into consideration the opinions of their constituency, thereby encouraging involvement in decision-making similar to direct or participatory democracy.

Unlike traditional decision support systems that are based on the experiences of small groups of individuals, we utilize elicited feedback from a large population to infer representative social decision models that maintain the diversity of the population, yet enable communities of fundamental interests and commitments to emerge. Collaborative filtering methods open up the decision process to large groups, but their predictions are based on shallow beliefs with no context or reasoning to back them up. Our approach builds models from elicited feedback that capture the contextual complexities involved in decision-making, such as cause and effect and

environmental uncertainties. These models of collective cognition will help guide decisions and increase understanding of human systems. We believe that increasing diversity in the decision-making process will increase the quality of solutions, reduce groupthink, bias and other issues that typically arise in group decision-making.

In this paper we discuss our progress in developing our social decision-making approach. We describe our Bayesian belief aggregation technique and how we apply it to probabilistic decision models. We apply our approach to the issue of stem cell research using opinions collected from an online resource for eliciting human intelligence called Mechanical Turk (mturk.com). We show how clustering across beliefs enables groups to emerge with opposing beliefs and goals. We contrast our approach with one that forces a cooperative situation by merging all beliefs into a single consensus. We show that our approach enables one to more fully understand the situation including the representative groups that emerge and enables us to utilize game and decision theoretic techniques to make the appropriate decision.

This paper is organized as follows. Section III describes many research areas that influence our social decision modeling approach. Section IV describes our primary research area of Bayesian belief aggregation in more detail. Section VI discusses how we apply our belief aggregation approach to decision models. Section VII shows describes how we enable the competitive aspects in social decision problems to emerge. Finally, section VIII discusses our results of applying our approach to stem cell research issues.

### III. RELATED WORK

#### A. *Sociology*

Social decision-making is an area of research that addresses the manner that societies of individuals (human and non-human) collectively make decisions [1]. According to social scientist Margaret Gilbert, social action is an action by an individual motivated by another individual [2]. Social scientists diverge on the concept of collective beliefs and collective behavior, some indicating that individuals only act in self-interest [3]. Others state that the concept of belief requires a single “mind” and that social groups do not have “minds” [2]. Some social scientists, such as Emile Durkheim, believe that social beliefs can exist outside of individual consciousness, indicating that groups are the synthesis of human beings and their collective beliefs. According to Durkheim, “when and only when they have formed a social group a certain special type of phenomenon can occur among them.” [2] This belief evokes the concept

of emergence, in which complex behaviors emerge from the interaction of many simple parts [4].

Some structures in place in society attempt to simplify social decision-making. For example, governance enables decision-making that affects many individuals. Different forms of government systems consider the individuals beliefs and preferences to varying degrees. Democracy is an attempt to empower the individual such that each has a say in the issues that affect everyone, however direct democracy (in which everyone votes on all issues) does not typically exist in practice in large groups. In modern democratic governments individuals remain removed from most decisions that affect them and instead put their trust in elected officials. We will attempt to address some of these shortcomings of democracy in our research.

#### B. *Computational Models*

Technological progress has been made in decision-making and problem-solving using behavior and computational models. For instance, game theory is closely related to decision theory and provides mathematical models of the behavior of competing and cooperative individuals. Game theory is based on comparing the utility (or value) of an outcome for the players of a game [5]. The most competitive situation, called a zero-sum game, optimizes the utility for one player at the remaining players expense. A very different situation, called Nash equilibrium, occurs when no individual can stand to improve his situation by changing strategies assuming the other players maintain their current strategies. Another cooperative situation called Pareto optimality occurs when no individual can do better without a different individual doing worse.

Decision models such as decision trees, Bayesian networks and influence diagrams also attempt to enumerate the decisions, potential outcomes and side effects in a situation [6]. These models also incorporate uncertainty, thereby improving their accuracy in real-world situations. Bayesian networks contain nodes representing variables with probabilistic outcomes and edges representing dependencies between the variables. A directed edge from one node (a parent) to another (a child) indicates that there is a causal relationship such that the parent has some effect on the child. Influence diagrams are an extension of Bayesian networks that also include decision nodes representing decisions and their options as well as utility nodes representing the value (positive or negative) of an outcome. Influence diagrams allow one to determine the decision option that will maximize utility given the uncertainties in the model.

All of the aforementioned computational models suffer from a significant shortcoming when considering decisions that affect many individuals. They assume that beliefs, utilities and other model components represent either a single individual or a group of individuals who have reached a consensus. Belief aggregation [8], [9] and topological fusion [10] are techniques used to form probabilistic models from the beliefs of multiple individuals. These approaches utilize opinion pool functions to generate a single probability distribution typically a weighted average from the supplied beliefs. The problem is that as the size of groups involved in a decision grow, their diversity of opinion and preferences will also grow. Any attempt to find or force a consensus on divergent beliefs is prone to failure [8] and will not form a realistic representation of a population. Instead of forcing or requiring consensus, our proposed approach forms computational models that maintain the diversity of the sources.

### C. Collaborative Filtering

Collaborative filtering is the process of making predictions from many individual opinions on a given topic. Collaborative filtering algorithms are typically used to make a recommendation to individuals by suggesting other items that individuals with similar interests have enjoyed. Web retailers such as Netflix use collaborative filtering in this manner based on ratings supplied by customers and purchase history. Recent techniques include matrix factorization such as singular value decomposition (SVD) and combining a number of algorithms into a hybrid approach. These techniques are used by one of the leaders of the Netflix prize competition [11].

While our approach is similar to collaborative filtering in that a prediction is based on community input versus singular input, we move beyond existing techniques by building entire models from community input instead of simply providing a prediction on one item or one event. We elicit opinions that ask contributors to consider the reasoning behind their beliefs, not just isolated beliefs without context. This enables predictions at a much finer level of granularity, and allows us to build complex causal and contextual models that can provide recommendations as well as help decision-makers understand the possible outcomes as well as the factors that contribute to the outcomes.

### D. Agent Behavior Models

Countless agent behavior models exist that attempt to simulate human behavior using software agents. Many researchers develop cognitive architectures that generalize the human concepts of observation, reasoning,

memory, learning, communicating and acting [12]. These agents are then used in simulations of real-world situations, for instance a simulation attempting to answer why the Anasazi abandoned their homeland in the 1300s [13]. Another relevant agent-based model was Robert Axelrod's Culture Model that attempted to model beliefs and attitudes in a population over time [14]. Agent-based research has made considerable advances for autonomous computing and simulation; however they remain an oversimplification of the human psyche and collective human behavior. Instead of simulation, we develop models from direct human input that will facilitate actual collective behavior in human systems.

## IV. 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 [6]. 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* [15]. 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 [8].

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.

## V. 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. A “belief” can represent either an a priori probability estimate for a variable in the Bayesian network, or a conditional probability estimate for a single parent-child relationship. Instead of computing a single consensus value (or average) to aggregate the potentially divergent probability estimates for a belief, 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. Our approach is described in greater detail in [9].

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 subsets of beliefs. Our approach finds the “best” set of subgraphs in the network and clusters within the subgraphs. We use Kullback-Leibler divergence to compare the distribution of the clusters within the subsets of beliefs. A subset with a higher K-L divergence has a better separation between clusters. Clustering across subsets of beliefs reduces the complexity versus clustering for each individual belief, yet allows more specialized groups to emerge.

## VI. APPLICATION TO DECISION MODELS

We extend the concept developed in [16] 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.

In some cases a social decision problem will be strategic, such as the example described in [16]. In a strategic game, one player selects a strategy, or decision, based on the strategy he expects the other player to use. In this case, each player chooses an individualized strategy. While we will eventually apply our approach to strategic games, this paper describes a situation in which there is only one decision to be made, but the same decision effects everyone. In this situation the strategy is deciding which decision option would be most appropriate given all players’ goals.

The MAID in [16] assumes that an agent represents a single entity. In our case, we consider the aggregated beliefs and utilities of potentially many individuals. Our initial experiment described in this paper clusters across all beliefs in the decision network such that there are multiple “agents” that each represent a subset of individuals for all beliefs in the network. Our continued research will apply our approach described in Section V that clusters across subsets of beliefs. For decision networks it may be the case that conflicts emerge as to which decision is the best option for an individual, based on the clusters the agent belongs to. For example, this may occur if the individual has conflicting beliefs. We will address this interesting situation in future research.

## VII. GAME THEORETIC ANALYSIS

Using an opinion pool function that finds a single consensus from all beliefs, the result of inference on a decision network would be a single utility estimate for each decision option. The decision option with the highest utility would theoretically represent the “best” decision. However, when we allow the divergent groups to emerge, we instead have a utility for each group and we may discover that the “best” decision is actually undesirable for a particular group. Utilizing an opinion pool that finds a single consensus results in a “forced”

cooperative situation, whereas when we allow the divergent groups to emerge we see that the situation may in fact be competitive. This distinction may be important for decision-making because it enables them to visualize the groups that will be negatively impacted by what could initially appear to be the best decision.

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. In this paper we show that our clustering approach enables us to analyze whether a solution is Pareto optimal, meaning that no other solution would increase the utility for one group without lowering the utility for another. We also determine a “satisficing” solution that attempts to reduce the loss for all groups.

### VIII. A DECISION MODEL FOR STEM CELL RESEARCH

#### A. Decision Model and Data

In this paper we apply our approach to a topical issue that is currently being addressed by policy-makers. The issue of stem cell research is particularly appropriate because of its polarizing effect on individuals with diverse backgrounds and motivations. Our goal was to determine whether the government should fund embryonic stem cell research, adult stem cell research, both or neither given opinions collected from an online survey. Our dataset is composed of opinions from 293 people using Mechanical Turk (mturk.com), an online source for low-cost human labor, and Survey Monkey (surveymonkey.com). Fig. 7 shows the questions and their options that composed the survey. The questions were used to form the influence diagram in Fig. 1. The Key in Fig. 7 indicates the corresponding node in the influence diagram. The rectangle node at the top represents the decision options. The oval nodes are *chance* nodes, representing a random (probabilistic) variable. The diamond nodes are *utility* nodes representing the value of an outcome. A positive value is considered “good” while a negative value is undesirable.

#### B. Inference

Each belief option in the survey was transformed into a numeric value (shown next to each option in Fig. 1) such that an overall numeric utility for the network could be computed. The utility was computed by selecting a decision option and then summing the inferred utility value of each utility node. Inference on the network

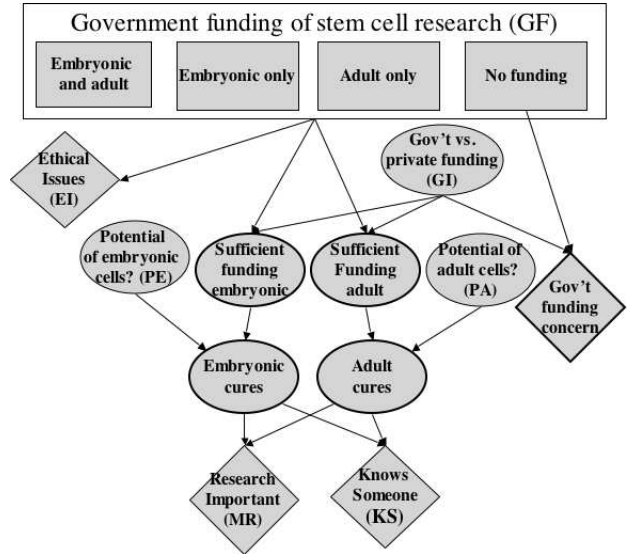


Fig. 1. An influence diagram for the stem cell research issue.

determines the value of each utility node based on which decision option was selected and the posterior probability of any chance nodes leading into the utility node. The posterior values of the nodes outlined with a heavier line were determined from their parent nodes. The a priori values of the parent chance nodes were derived from the user beliefs.

Appropriate utility values were determined from experimentation. We ran the inference algorithm on each individual’s supplied beliefs given the values we had selected for utilities. If the decision option that the individual had selected matched our highest utility option, we considered this a *correct* prediction. We attempted a number of combinations of utility values until we achieved the highest number of correct predictions. The remaining *incorrect* predictions corresponded to individuals who had selected decision options that did not seem to fit their beliefs and utilities. While some mispredicted answers may be due to misunderstanding of the questions, in some cases these situations are opportunities to guide the user to a decision that increases the user’s utility.

#### C. Comparing single consensus versus groups of consensus results

Fig. 2 shows the probability distribution for each belief given all individuals surveyed. Each cluster of bars represents one belief and its options. For instance, the first belief labeled with **GF** indicates the preferred options for the question “Should the government fund stem cell research?”. The first bar to the left is the

probability for decision option  $a$  in Fig. 7. The second bar corresponds to option  $b$  and so on.

In order to determine the “best” decision option, the overall utility of the network is computed for each option. We first computed the overall utility of the network using a single consensus value for each belief. In other words, the value for each belief was a linear average of all the users’ inputs. This approach resulted in the following utilities, indicating that “Fund both embryonic and adult stem cell research” was the best decision option for the government.

Decision option	Utility
<i>Fund embryonic and adult stem cell research</i>	16.3
Fund embryonic stem cell research only	6.1
Fund adult stem cell research only	13.8
Do not fund any research	3.8

While this approach determines the “best” decision option given the averages of the beliefs, it does not represent a deep understanding of the underlying opinions. In fact, there is a significant subset of the population (8%) who believe that the government should not fund any research according to Fig. 2. While a majority of the surveyed population may believe that the government should fund both embryonic and adult stem cell research, policy-makers may also be interested in who is likely to oppose such measures. They may want to include in any new policy a compromise to deal with a potentially vocal minority.

Our second experiment first finds a suitable clustering of the population based on their supplied beliefs before computing the overall utility. We tested two different clustering techniques: a graph-based normalized cut algorithm [9], [17] and Weka’s simple k-means clusterer [18]. The results were slightly better with the k-means clusterer based on the computed entropy and standard deviation. The results shown in this paper utilized the k-means clusterer. Figs. 3 through 6 show the probability distributions of each cluster, each of which has very distinct characteristics based on its members’ beliefs and concerns. Cluster 1 strongly supports embryonic research and has a high personal stake in the outcome based on knowing someone who could potentially benefit from research. Cluster 2 overwhelmingly supports private funding instead of government funding for research. This group has a lower faith in the potential of stem cell research in general. Cluster 3 indicates that the government should only fund adult stem cell research and has strong ethical concerns about embryonic stem cell research. Cluster 4 supports embryonic research but has a lower personal stake than the first cluster.

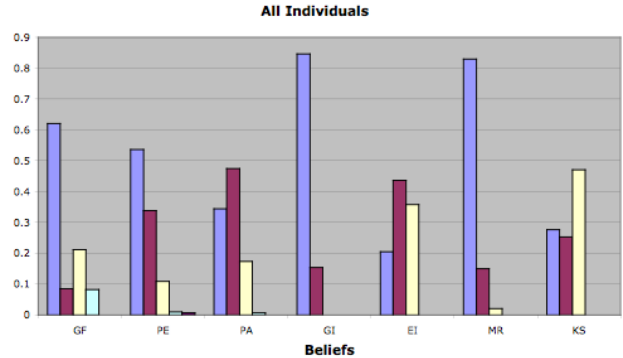


Fig. 2. The probability distribution for each belief given all individuals surveyed. The beliefs are shown in the x axis, with each color indicating a different option. The abbreviations under each cluster of bars correspond to the key in Fig. 7

Fig. 8 shows the overall utility for each cluster given each decision option. We can analyze these results to determine several interesting characteristics that can help a decision-maker choose the most appropriate option. These characteristics are not distinguishable using a single consensus approach. First of all we can determine the Pareto optimal solutions in which no group can do better without another group doing worse. In this case, all solutions are Pareto optimal except for funding embryonic research only, an option that all groups can do better than. We also learn from these results the optimal solution for each group as well as the worst case scenario for each group. A minimax solution would be one in which, given a strategic game, a player chooses a strategy that minimizes her maximum possible loss. While our decision problem is not strategic, we can still choose an option that is not the worst case scenario for anyone. This solution would be to fund adult stem cell research only. This may be an appropriate solution if a policy-maker wishes to make a compromise. We call this a “satisficing” solution. In this case the satisficing solution is not the same as the optimal solution discovered using the single consensus approach.

## IX. CONCLUSIONS AND FUTURE WORK

In this paper we demonstrated our initial results using our social decision-making approach. This approach forms decision models from the input of many individuals while maintaining the divergence in belief and values represented by a diverse population. Our approach causes the competitive aspects in a decision problem to emerge and enables one to utilize game and decision theoretic analyses to determine an appropriate solution. We see the potential of our approach to assist decision- and policy-

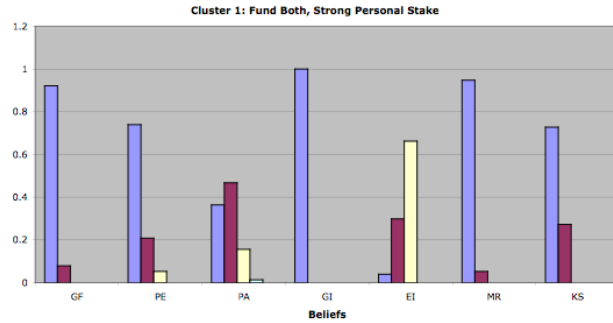


Fig. 3. The probability distribution for the first cluster, indicating strong support for embryonic research and a high personal stake in the outcome.




Fig. 6. The probability distribution for the fourth cluster, indicating that the government should fund both types of research, but with a lower personal stake than in cluster 1.




Fig. 4. The probability distribution for the second cluster, which overwhelmingly supports private funding for research.

makers and increase the involvement of individuals who have a stake in the outcome of a situation.

Our continued research will apply our approach to strategic games, in which multiple decisions are made by a number of individuals. We will also apply our approach that clusters across subsets of beliefs in decision models to improve the specificity of the clusters and help discover individuals with complex or conflicting beliefs.




Fig. 5. The probability distribution for the third cluster, indicating that the government should only fund adult stem cell research.

## REFERENCES

- [1] A. G. Sanfey, "Social decision-making: insights from game theory and neuroscience," *Science*, vol. 318, no. 5850, pp. 598–602, 2007.
- [2] M. Gilbert, *On Social Facts*. Princeton University Press, 1992.
- [3] M. Olson, *Logic of Collective Action*. Harvard University Press, 1971.
- [4] J. Fromm, "Ten questions about emergence," 2005. [Online]. Available: <http://arxiv.org/abs/nlin/0509049v1>
- [5] K. Leyton-Brown and Y. Shoham, *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, 2008.
- [6] K. Murphy, "A brief introduction to graphical models and bayesian networks," 1998. [Online]. Available: <http://www.cs.ubc.ca/~murphyk/Bayes/bnintro.html>
- [7] D. M. Pennock and M. P. Wellman, "Graphical representations of consensus belief," in *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 1999.
- [8] K. A. Greene and G. F. Luger, "Agreeing to disagree: Leveraging consensus and divergence in bayesian belief aggregation," in *Papers from the AAAI Spring Symposium, Technosocial Predictive Analytics Workshop*, 2008, pp. 49–54.
- [9] I. Matzkevich and B. Abramson, "The topological fusion of bayes nets," in *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*, 1992.
- [10] R. Bell, Y. Koren, and C. Volinsky, "The bellkor 2008 solution to the netflix prize. document generated for the 'netflix prize' competition," 2008. [Online]. Available: <http://research.att.com/~volinsky/netflix/Bellkor2008.pdf>
- [11] R. Wray, R. Chong, and J. Phillips, "A survey of cognitive and agent architectures," 2009. [Online]. Available: <http://ai.eecs.umich.edu/cogarch0/>
- [12] R. Goldstone and M. Janssen, "Computational models of collective behavior," *Trends in Cognitive Sciences*, vol. 9, no. 9, pp. 424–430, 2005.
- [13] R. Axelrod, "The dissemination of culture: a model with local convergence and global polarization," *The Journal of Conflict Resolution*, vol. 41, no. 2, pp. 203–226, 1997.
- [14] R. Howard and J. Matheson, Eds., *Readings on the Principles and Applications of Decision Analysis*. Strategic Decisions Group, 1984, vol. 2, pp. 6–16.
- [15] D. Koller and B. Milch, "Multi-agent influence diagrams for representing and solving games," in *IJCAI*, 2001, pp. 1027–1036.
- [16] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, 2000.
- [17] I. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2005.

Key	Question	Options
GF	Should the government fund stem cell research?	a. Yes, the government should fund embryonic and adult research b. The government should only fund embryonic stem cell research c. The government should only fund adult stem cell research d. The government should not fund any stem cell research
PE	What is the potential for embryonic stem cell research to cure diseases and provide new methods to test cures?	a. The potential is very high (90%) b. There is some potential (75%) c. It is not clear yet what the potential is (50%) d. It is unlikely to provide cures (25%) e. It is very unlikely to provide cures (10%)
PA	What is the potential for adult stem cell research to cure diseases and provide new methods to test cures?	a. The potential is very high (90%) b. There is some potential (75%) c. It is not clear yet what the potential is (50%) d. It is unlikely to provide cures (25%) e. It is very unlikely to provide cures (10%)
GI	Would government funding of stem cell research improve its potential?	a. Yes, government funding is important to advance research (99%) b. No, private funding is sufficient to advance research (1%)
EI	Are you concerned about the ethical issues of using embryonic stem cells?	a. I am very concerned. The potential does not outweigh the ethical issues. (-20) b. I am a little bit concerned. I think there is potential, but some issues need to be addressed. (-10) c. I am not very concerned. I think the potential outweighs any ethical issues. (0)
MR	How important is advancing medical research and curing diseases in general to you?	a. Very important. (+15) b. Somewhat important. (+10) c. Not very important. (0)
KS	Do you know someone who could potentially benefit from stem cell research?	a. Yes, myself or a loved one (+15) b. Yes, a friend or associate (+10) c. Not personally(0)

Fig. 7. Survey questions and options given to Mechanical Turk workers. Each question represents one belief. The Key shows the abbreviation for the belief.

Decision	Cluster				Pareto Optimality
	Cluster 1	Cluster 2	Cluster 3	Cluster 4	
<i>Fund Both</i>	38.64	-19.42	14.81	16.74	Optimal
<i>Fund Embryo only</i>	18.82	-9.29	-0.15	5.75	Not optimal
<i>Fund adult only</i>	19.83	3.35	14.97	11.17	Optimal
<i>Fund neither</i>	0.01	19.45	0.01	0.17	Optimal
<b>Characteristics</b>					
<i>Size</i>	77	44	56	116	
<i>Percent</i>	26%	15%	19%	40%	
<i>Optimal solution</i>	Fund both	Fund neither	Adult only	Fund both	
<i>Worst solution</i>	Fund neither	Fund both	Embryo only	Fund neither	

Fig. 8. Overall utility results using belief clusters and some characteristics of the clusters .