Proceedings of the 19th Conference on Behavior Representation in Modeling and Simulation, Charleston, SC, 21 - 24 March 2010

Data-Driven Coherence Models

Peter Danenberg
Computer Science Department
University of Southern California
941 Bloom Walk
Los Angeles, CA 90089-0781
peterchd@usc.edu

Stacy Marsella
Institute for Creative Technologies
13274 Fiji Way
Marina del Rey, CA 90292
marsella@ict.usc.edu

Keywords: Thagard, coherence, Iraq, data-driven

Abstract

How groups maintain and revise their beliefs and attitudes in the face of new information is a basic research question in human social behavior and communications, as well as having a range of applications in crafting effective communications in such areas as health interventions, political campaigns and advertising. In this paper, we argue for a coherence based approach for modeling group belief revision processes and as a framework for studying belief and attitude change. Coherence models have a rich history of applicability in the psychological sciences where they have been used to explain a range of belief maintenance processes in the individual. Given that processes of social comparison and pressure can homogenize a cohesive group's beliefs, we argue in this paper to extend the application of coherence models to modeling group belief systems. Additionally, we address challenges in constructing and using such belief models. Typically, creating accurate models of either an individual or group's beliefs requires the painstaking engagement of domain experts. We present and demonstrate a method for producing them from data and exploring potential vectors of attitude change in their subpopulations.

1 Introduction

1.1 Thagard's coherence

Coherence has been proposed as a general cognitive mechanism by which, for instance, a person forms explanations (Bonjour, 1976), integrates information to form impressions of others (Rawls, 1974-5); and resolves cognitive dissonances between beliefs and behavior (Festinger, 1954).

Coherence is part of a rich history of philosophical debate. Bosanquet argues (Bosanquet, 1912, p. 340) that it stretches back to Plato's theory of forms; where a set of N manifestations asymptotically coheres towards its universal form; and even in Hegel's dialectic, where disparate or indeed antithetical elements cohere in the process of "sublation."

By contrast, Aristotle's critique of Platonic realism lays the foundations of empiricism; and the Platonic-Aristotelean breach eventually leads to the foundationalism vs. coherentism debate (BonJour, 1985; Moser, 1988a; BonJour, 1988; Moser, 1988b). Whereas the former argues that epistemological justification "requires a non-propositional basis in the contents of experience;" (Moser, 1988c) the latter maintains that "beliefs are justified by being inferentially related to other beliefs in the overall context of a coherent system." (Bonjour, 1976)

Thagard establishes his system of "coherence as constraint satisfaction", we argue, by drawing from coherentist and foundationalist models of justification. One determines, for instance, the justification of a belief *vis-à-vis* its explanatory corroboration by other beliefs in its system with which it's associated (Thagard and Verbeugt, 1998, p. 155); but surprisingly, perhaps, Thagard gives priority to beliefs from observation (Thagard and Verbeugt, 1998, p. 157). BonJour calls this 'weak foundationalism,' whereby the "initial modicum of justification [for empirical beliefs] must be augmented by a further appeal to coherence before knowledge is achieved." (Bonjour, 1976, p. 284)

1.2 Attitude change

We'd like to address the problem of attitude change, proposing a practical method for identifying potential vectors of communicative hegemony; of interest in health intervention, political campaigns and marketing propaganda.

Finding the right communication to persuade someone, however, often hinges on tailoring it to their attitudes and beliefs. This suggests a dual-pronged approach of exploring alternative messages even while tailoring them to potentially receptive subgroups. Such a dual-pronged approach requires searching two spaces simultaneously: the space of possible message contents and the space of possible subgroups to which the message will be conveyed. To avoid searching the Cartesian product of message-subgroups, we can identify subgroups based on whether they share a common coherence model that is amenable to change and then use that model to suggest approaches to attitude change (see section 4.1, "Perturbation"). We take for granted, however, that coherence mechanisms provide a way to optimize messages for a given subgroup.

Coherence models in psychology, however, have largely been seen as cognitive mechanisms operating within the individual. The strong view of our approach is to argue that the coherence mechanisms also operate in group attitude change; nevertheless, a weaker view may be sufficient: e.g. finding a stereotypical, average individual of a group for which the message works.

The argument for extending coherence to modeling groups follows from several classic theories in social psychology. Most notably, Festinger's work on social comparison theory (Festinger, 1954, p. 125) that argues that individuals have a need to assess their beliefs by comparison with others. Festinger's work suggests that groups strive for a quiescent homogenization of opinion; and to that end tend to exclude discrepant members, pressure non-discrepant ones towards uniformity As a result, groups evince a principle of spontaneous self-cohesion not unlike the reduction of cognitive dissonance in individuals. Similar views can be be seen in more recent theories as social appraisal theory.

We argue, therefore, that persuasive messages targeted at groups will demonstrate a similar attitude-mutating effect across its members.

Thagard's doctrine of coherence as constraint satisfaction provides our point of departure (Thagard and Verbeugt, 1998); whose models, however, are laboriously forged by domain experts relying on intuition. Our counterproposal, therefore, is a data-driven approach whose process is threefold:

- inducing structural models from survey data;
- "drilling down" into the beliefs of subgroups exposed by the data;
- perturbation of the subgroup-models to expose mutable attitudes as potential targets of persuasion.

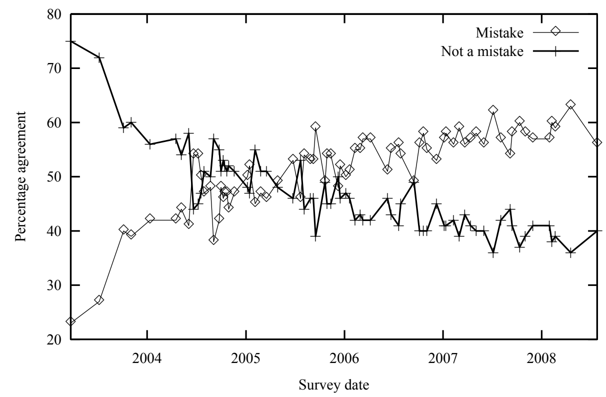

Figure 1: Gallup survey, "Do you think the United States made a mistake in sending troops to Iraq?"

By way of case study, we apply our method to public opinion around the Iraq War.

2 Motivating Example: Iraq

The Iraq war was a highly polarizing event. A January 2007 poll showed that roughly three-quarters of the world's population disapproved of how the U.S. policy on Iraq (BBC World Service, 2007); American opinion had a relatively constant, even bipartition from 2004 until 2006 (Gallup, Inc., 2008), when opposition to the Iraq War began to increase by a widening margin (figure 1).

The Iraq War struck us as a potentially fertile ground for studying attitude change, given the volatile and strong, even radicalizing, nature of people's opinion on the matter; and, indeed, motivating people to provide data was relatively simple (see section 5).

3 Coherence Model

Our working model of attitude stability and mutation is based on Thagard's formalization of coherence as constraint satisfaction (Thagard and Verbeugt, 1998): *videlicet*, the partitioning of a system of propositions E into disjoint subsets A and R; corresponding to accepted and rejected propositions, respectively. The propositions themselves,

$$\{e_1, e_2, \ldots, e_n\}$$

are subject to the weighted constraints

$$\{(e_{i1}, e_{j1}), (e_{i2}, e_{j2}), \ldots, (e_{in}, e_{jn})\}$$

such that

$$((e_i, e_j) \in C+ \rightarrow e_i \in A \leftrightarrow e_j \in A) \wedge$$
$$((e_i, e_j) \in C- \rightarrow e_i \in A \leftrightarrow e_j \in R)$$

where $C+$ and $C-$ are sets of positive and negative constraints. The coherence problem, then, becomes the maximization of W; *id est*, the sum of all satisfied constraints' weights.

Although the coherence problem is NP-complete (Thagard and Verbeugt, 1998, page 2), there exist a number of approximating algorithms; from which we chose the connectionist for its natural affinity to coherence problems (see section 3.1) and general applicability.

The connectionist model has been variously described as:

- minimizing the "energy" of a system though gradient descent (Sejnowski, 1986; Hopfield, 1982);
- maximizing the "harmony" of a system (Smolenksy, 1986);
- maximizing the "goodness-of-fit" of a system's contraints, such that

  $$G(t) = \sum_i \sum_j w_{ij} a_i(t) a_j(t) + \sum_i input_i(t) a_i(t)$$

  where w corresponds to the weight of a constraint, a to a node's activation, and $input$ to an imposed bias (Rumelhart and McClelland, 1986).

3.1 Goodness-of-fit

Thagard characterizes coherence as constraint satisfaction by abstracting upon Rumelhart's goodness-of-fit (Thagard and Verbeugt, 1998, page 10); and generalizes away, in particular, the latter's adherence to neural networks. Armed with his abstracting coherence, Thagard is able to reformulate classic problems across several areas of research, including:

psychology: cognitive dissonance (Schultz and Lepper, 1996), interpersonal relations (Read and Marcus-Newhall, 1993);

politics: deliberate democracy (Arrow, 1963; Black, 1998);

ethics: reflective equilibrium (Daniels, 1979; Reuzel et al., 2001).

Spellman, et al. (Spellman et al., 1993) adapt Thagard's coherence model to simulate attitudinal shifts during the First Gulf War; which adaptation they characterize as "dissonance reduction." Proceeding from a hand-crafted network of attitudinal relations, they capture the maintenance of cognitive consistency across attitude-shifting events; which corroborates survey data they gathered and independently analysed.

Going beyond Thagard, we've developed a technique of perturbation (*vide* section 4.1) or *subjunctive* constraint satisfaction; whereby we determine, for any given target node, its prime hegemons.

Coherence models are typically hand-crafted by researchers and other domain experts (Thagard, 2003); requiring not only extensive knowledge but also subject to gaps in knowledge and biases. What follows is a method to create coherence models directly from data.

4 Data-driven Model Construction

Spirtes *et al.* developed a search algorithm for discovering causal structures from data, which they called the "PC algorithm" (Spirtes et al., 2000, p. 84). It starts by forming a complete undirected graph (whose vertices correspond to random variables), deleting conditional independencies and orienting the remaining links according to Pearl's IC algorithm (Pearl, 2000, page 50).

Assuming that the functions Adjacent(G, i, j) and Adjacencies(G, i) have been defined, which return whether i and j are adjacent in graph G, and all the vertices adjacent to i in G, respectively.

The SGS algorithm, predecessor to PC, had an expected running time of $\Omega(k^n)$; which PC has improved to $O(n^k)$ by testing fewer d-seperations in the case of sparse DAGs. (That a given DAG be sparse is often a reasonable assumption (Kalisch and Bühlmann, 2007, page 2).) PC works, namely, by incrementally removing conditional independencies of order $0 \leq k \leq n$; where n is the cardinality of the largest set **k** d-separating some nodes i and j. Its performance is therefore inversely proportional to the connectedness of a given graph.

4.1 Perturbation

The skeleton C returned by PC-Algorithm is a coherence-like model suitable for exploration by perturbation.

For a given target node t among nodes $\{v_1, v_2, \ldots, v_n\} = V$ in a coherence network, perturbation individually sets the activation of $v_i \in V$ to min-activation or max-activation, runs the connectionist algorithm, notes the divergence of t's activation, and performs a partial ordering of V for each t by $\max(|\Delta_{ti\min}|, |\Delta_{ti\max}|)$ into non-, weak- and strong-hegemons.

4.2 Method

4.2.1 pcalg

Data is collected, stored and imported into R; the `pcalg` (Kalisch and Maechler) package is then used to create an apposite skeletal UDAG, and specialize this UDAG into one of an equivalence class of underlying DAGs.

4.2.2 Influence

The underlying DAGs are then imported into Influence, a reimplementation of Thagard's ECHO by Danenberg, et al.; via one of two methods:

1. a Scheme-to-Java bridge implemented in SISC;
2. a custom R server on an arbitrary machine.

Once in Influence, one can create arbitrary cross-sections of the data by subsetting on demographics or response; and from this cross-section, recreate the graph structure (including node activations and internodal relationships).

Next, the graphs of sufficiently interesting subpopulations can be perturbed and compared; and their structural differences reasoned upon (see section 5).

5 Experiment

For the survey instrument, we assembled twenty-eight items on a five-point Likert scale; with a demographic section covering education, ethnicity, income and party affiliation. As of this writing, the survey is still available on-line (Danenberg, 2007).

We solicited for subjects on Google AdWords (Google, Inc., 2008) from March 27–29, 2007 under the slogan: "We need your opinion on Iraq. Take our Iraq War survey!" The cost of the campaign was $\$1451.29$; and of the $473,685$ ad impressions, we had 627 visits; of those visits, 442 surveys were submitted; of those surveys, 98 were rejected for incompleteness: leaving 344 valid responses.

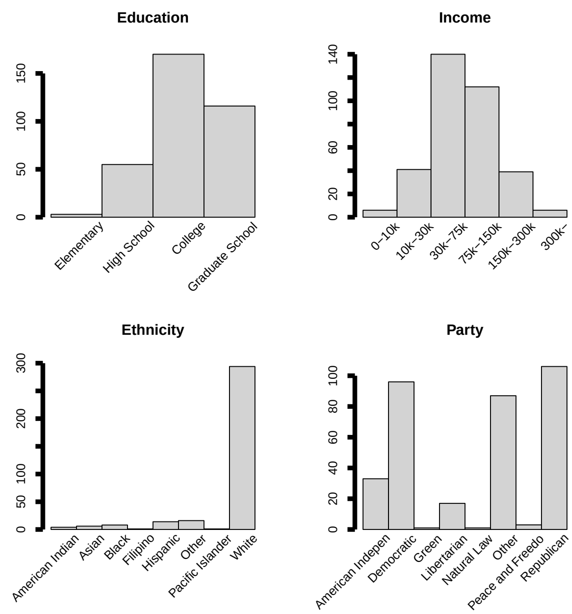

Figure 2: Histogram of respondents over education, ethnicity, income and party

Figure 2 summarizes the demographic data. Although education, $\chi^2(2, N = 341) = 468, p < 0.001$ (Stoops, 2004); ethnicity, $\chi^2(2, N = 322) = 76.4, p < 0.001$ (Survey, 2006); and income, $\chi^2(2, N = 187) = 95.3, p < 0.001$ (U.S. Census Bureau, 2006) defied the census; party affiliation, $\chi^2(2, N = 344) = 7.62, p < 0.05$ compared favorably with the latest Pew statistics (Pew Research Center, 2008), but that fewer Democrats filled out the survey than expected (table 1).

	Republicans	Democrats	Others	Total
Observed	106	96	142	344
Expected	96.3	120.4	127.3	344
Residuals	0.986	-2.224	1.305	0.67

Table 1: Observed vs. expected party affiliation

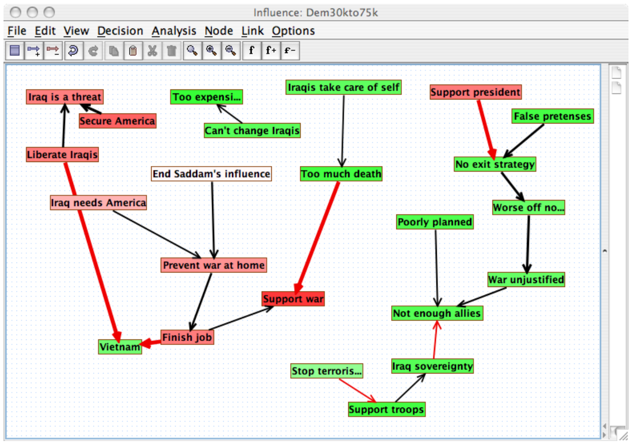

Figure 3: Poor Democrats

5.1 Model

Figures 3, 4, 5, 6, show the models gleaned from subsetting the data by income (poor/rich) and party (Democrat/Republican) after analysis; whose relative sparseness compared to the full model is proportional to their data-density.

5.2 Analysis

Table 2 summarizes the perturbation results on all four subgroups; a striking observation whereof is how class runs thicker than party: rich Democrats and Republicans are repulsed by the war's cost ("Too expensive"), while poor Democrats and Republicans are repulsed by its inhumanity ("Vietnam").

Almost universally, however (with the exception of rich Democrats, for whom we lack data), "Support the president" positively correlates with "Support war" (figures 3, 5, 6); even though Democrats and Republicans differ across party lines.

Amongst poor Democrats and poor Republicans (figures 3, 5), "Vietnam" appears to be associated with "Too much death;" we speculate the cause being that American casualities in Iraq are predominantly poor (Scotland, 2008).

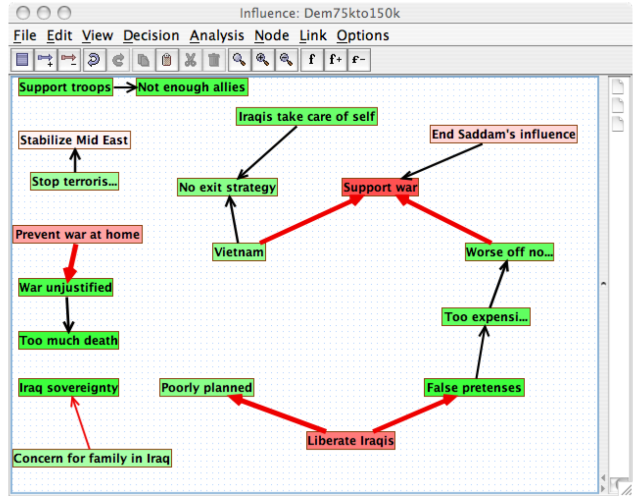

Figure 4: Rich Democrats

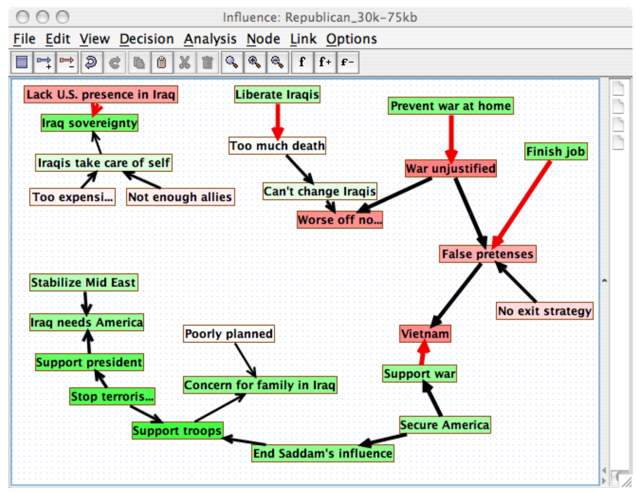

Figure 5: Poor Republicans

Rich Democrats and rich Republicans (figures 4, 6), on the other hand, demonstrate correlation between "No exit strategy" and "Too expensive;" could it be that they foresaw asymmetrical taxes on this liability (Montopoli, 2009)?

6 Conclusion

Creating coherence models by hand is an error-prone activity which beggars, furthermore, one's ingenuity; we present a method for creating models from data and identifying potential vectors of attitude change through perturbation.

	Republican		Democrat	
Influence	Poor	Rich	Poor	Rich
Positive				
Strong	Iraq needs America	n.d.[a]	n.d.	n.d.
Moderate	n.d.	Liberate Iraqis Stabilize Mid East Iraq needs America Concern for family in Iraq Prevent war at home Finish job Secure America	Secure America Finish job Stabilize Mid East Prevent war at home	n.d.
Weak	n.d.	n.d.	n.d.	n.d.
Negative				
Strong	Vietnam	n.d.	n.d.	n.d.
Moderate	n.d.	Worse off now War unjustified Poorly planned Vietnam No exit strategy Can't change Iraqis Not enough allies Too expensive Iraqis take care of self	Vietnam False pretenses No exit strategy	Too expensive Worse off now Not enough allies Can't change Iraqis
Weak	n.d.	n.d.	n.d.	n.d.

[a]No data

Table 2: Perturbation on "Support war" for poor/rich Democrats/Republicans

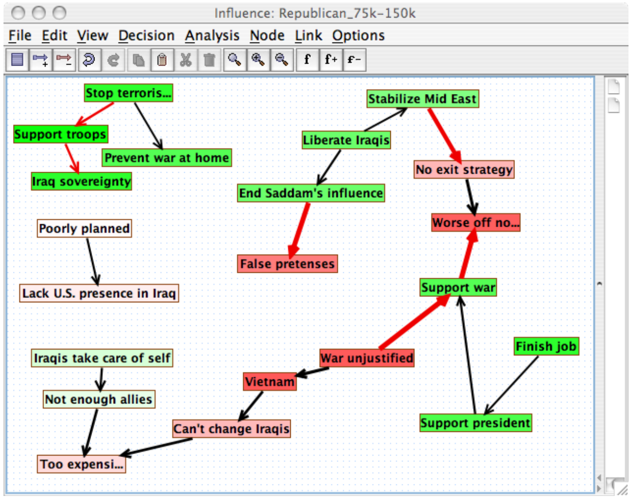

Figure 6: Rich Republicans

We'd like to test mutating craft of the thus prescribed vectors in a follow-up study, wherein appropriate or inappropriate messages preface the administration of the instrument and attitude deviation is tested against the null hypothesis.

We're also evaluating the utility of the method for marketing and political campaigns.

References

K. J. Arrow. *Social choice and individual values*. Wiley, 1963.

BBC World Service. World view of us role goes from bad to worse. `http://news.bbc.co.uk/1/shared/bsp/hi/pdfs/23_01_07_us_poll.pdf`, January 2007.

Duncan Black. *The Theory of Committees and Elections*. Kluwer, 1998.

Laurence Bonjour. The coherence theory of empirical knowledge. *Philosophical Studies*, 1976.

Laurence BonJour. *The structure of empirical knowledge*. Harvard, 1985.

Laurence BonJour. Reply to moser. *Analysis*, 48(4):164–5, October 1988.

Bernard Bosanquet. *The Principle of Individuality and Value*. Macmillan, 1912.

Peter Danenberg. Iraq war survey. `http://iraq-war-survey.com`, March 2007.

N. Daniels. Wide reflective equilibrium and theory acceptance in ethics. *Journal of Philosophy*, 76:256–82, 1979.

Leon Festinger. A theory of social comparison processes. *Human Relations*, 7:117–40, 1954.

Gallup, Inc. Gallup's pulse of democracy: The war in iraq. `http://www.gallup.com/poll/1633/Iraq.aspx`, 2008.

Google, Inc. Adwords. `https://adwords.google.com`, 2008.

J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 9:2554, 1982.

Markus Kalisch and Peter Bühlmann. Estimating high-dimensional directed acyclic graphs with the pc-algorithm. *Journal of Machine Learning Research*, 8:616–36, 2007.

Markus Kalisch and Martin Maechler. pcalg: Estimating the skeleton and equivalence class of a dag. `http://cran.r-project.org/web/packages/pcalg/index.html`.

Brian Montopoli. David obey calls for war tax on wealthy. `http://www.cbsnews.com/blogs/2009/11/23/politics/politicalhotsheet/entry5748813.shtml`, November 2009.

Paul K. Moser. Internalism and coherentism: a dilemma. *Analysis*, 48(4):161–3, October 1988a.

Paul K. Moser. How not to be a coherentist. *Analysis*, 48(4):166–7, October 1988b.

Paul K. Moser. The foundation of epistemological probability. *Erkenntis*, 28(2):231–51, 1988c.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

Pew Research Center. Fewer voters identify as republicans. `http://pewresearch.org/pubs/773/fewer-voters-identify-as-republicans`, 2008.

John Rawls. The independence of moral theory. *Proceedings and Addresses of the American Philosophical Association*, 48:5–22, 1974-5.

S. Read and A. Marcus-Newhall. The role of explanatory coherence in the construction of social explanations. *Journal of Personality ctnd Social Psychology*, 65:429–47, 1993.

R.P.B. Reuzel, G.J. van der Wilt, H.A.M.J. ne Have, and P.F de Vries Robbé. Interactive technology assessment and wide reflective equilibrium. *Journal of Medicine and Philosophy*, 26(3):245–61, 2001.

David E. Rumelhart and James L. McClelland, editors. *Parallel distributed processing: explorations in the microstructure of cognition*. MIT Press, 1986.

T. R. Schultz and M. R. Lepper. Cognitive dissonance reduction as constraint satisfaction. *Psychological Review*, 103:219–40, 1996.

Herald Scotland. Us death toll in iraq is 'mostly white and poor'. http://www.heraldscotland.com/us-death-toll-in-iraq-is-mostly-white-and-poor-1.877356, March 2008.

T. J. Sejnowski. *Parallel distributed processing: explorations in the microstructure of cognition*, volume 2, chapter 6, "Parallel distributed processing: explorations in the microstructure of cognition", pages 372–89. MIT Press, 1986.

P. Smolenksy. *Parallel distributed processing: explorations in the microstructure of cognition*, volume 2, chapter 7, "Neural and conceptual interpretation of PDP models". MIT Press, 1986.

Barbara A. Spellman, Jodie B. Ullman, and Keith J. Holyoak. A coherence model of cognitive consistency: Dynamics of attitude change during the persian gulf war. *Journal of Social Issues*, 1993.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2000.

Nicole Stoops. Educational attainment in the united states: 2003. Technical report, U.S. Census Bureau, June 2004.

American Community Survey. B02001. race - universe: Total population. Technical report, U.S. Census Bureau, 2006. URL {http://factfinder.census.gov/servlet/DTTable?_bm=y&-context=dt&-ds_name=ACS_2006_EST_G00_&-mt_name=ACS_2006_EST_G2000_B02001&-CONTEXT=dt&-tree_id=306&-redoLog=false&-all_geo_types=N&-geo_id=01000US&-currentselections=ACS_2006_EST_G2000_B02001&-search_results=01000US&-format=&-_lang=en}.

Paul Thagard. Why wasn't O.J. convicted? emotional coherence in legal inference. *Cognition and Emotion*, 17 (3):361–83, 2003.

Paul Thagard and Karsten Verbeugt. Coherence as constraint satisfaction. *Cognitive Science*, 1998.

U.S. Census Bureau. Current population survey, 2006 annual social and economic supplement. http://pubdb3.census.gov/macro/032006/perinc/new01_001.htm, 2006.

Author Biographies

PETER DANENBERG is a humanist-turned-hacker and graduate student at USC, specializing in cognitive modelling.

STACY MARSELLA is a research associate professor at the University of Southern California and an Associate Director of Research for Social Simulation at the Institute for Creative Technologies. He has worked for a number of years in the areas of planning, cognitive science, computer-assisted learning and simulation. He received his Ph.D. from Rutgers University.