Second International Symposium on Engineering Systems
MIT, Cambridge, Massachusetts, June 15-17, 2009

Decision-Making by Technical Expert Committees for Engineering Systems*

David A. Broniatowski[1]
Engineering Systems Division, Massachusetts Institute of Technology, 77 Massachusetts Ave. Room E38-546, Cambridge, MA, 02141

Abstract: *Engineering systems are complex by definition. Information about these systems must therefore be pooled from multiple technical experts. Understanding the social dynamics underlying multi-stakeholder decisions by committees of experts is therefore crucial to the nascent field of Engineering Systems. This research argues for the importance of the study of committee decision-making to the field of Engineering Systems, and provides a computational, quantitative approach for these analyses.*

Key words: Multiple stakeholders, Bayesian inference, committee decision-making, engineering systems

1. Introduction

Engineering systems necessarily involve multiple stakeholders. The behavior of a sufficiently complex engineered system is beyond the capacity of a single human mind to understand. Although a system architect can understand the gross behavior of a system at a high level of abstraction, interactions between seemingly unrelated low-level technical components are unlikely to be captured. Knowledge of and experience with the inner workings of these components is spread among expert specialists. Any large-scale engineered system must also receive the approval of several stakeholders, many of whom have differing requirements, and hence different perceptions, of the system and its functionality. Examples include design reviews that large-scale engineered systems must pass (consider, for example, the PDR and CDR cycles within the aerospace domain). These approval activities bring additional expertise to bear on improving the ultimate design. Information about the engineering system is therefore distributed among many different experts.

---

* This paper synthesizes and builds upon earlier versions of this research, presented at the following conferences:

[1] D.A. Broniatowski, "A Method for Generating Social Networks from Meeting Transcripts," *International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Modeling Intercultural Collaboration and Negotiation (MICON)*, Pasadena, CA: Springer-Verlag, 2009.
[2] D.A. Broniatowski, M. Yang, J.F. Coughlin, and C.L. Magee, "Bayesian Analysis of Decision Making in Technical Expert Committees," Conference on Systems Engineering Research, Loughborough, UK: 2009, p. 10.
[3] D.A. Broniatowski, C.L. Magee, J.F. Coughlin, and M. Yang, "Quantitative Analysis of Group Decision Making for Complex Engineered Systems," IEEE Systems Conference, Vancouver, British Columbia, Canada: 2009.

[1] Graduate Research Assistant, Engineering Systems Division, Massachusetts Institute of Technology, 77 Massachusetts Ave. Room E38-546, Cambridge, MA, 02141

Different experts, having been trained in different areas or components, will tend to pay attention to those elements of the system that they find consistent with their professional training – i.e., cognitively salient [4]. The mechanisms by which this training is achieved require acculturation within specific professional specialties, including learning that professional institution's language and jargon [5]. By institution, we mean a set of social norms to which a particular community adheres. This leads to a situation wherein individual experts develop different views of the system. In such cases, the system becomes a boundary object (cf. [6]), knowledge about which must be jointly constructed by the experts in question.

The committee is a common means by which experts pool their knowledge in an attempt to reach a consensus decision about a complex system or process. A successful committee will be able to integrate the disparate knowledge and viewpoints of its members so as to make a decision that solves the problem at hand. On the other hand, an unsuccessful committee can fail for many reasons – these include, but are not limited to, the absence of relevant technical expertise; the inability of committee members to communicate across disciplinary boundaries; and personality conflicts (see, e.g., [7] for an example of these challenges in the early FDA drug and medical device approval committees). Evaluating committee decision processes requires a means of understanding the social dynamics among members, and how these dynamics interact with the technical specifics of the system in question.

2. Literature Review

Although the dynamics of committee decision-making have been studied extensively from the perspective of social-choice theory, this body of literature proceeds from the assumptions that preferences are fixed in advance and that decisions are made on purely strategic grounds. Therefore, the differences between actors are modelled by different individual preference profiles [8]. Similar work in economics has differentiated between actors by "type" (i.e., expert or non-expert), with the intention of identifying an optimal allocation of experts [9]. These models have contributed much to our understanding of the rational-choice components of group decision processes. The research presented in this paper examines the social antecedents of preference formation. If preferences have been determined successfully, it is conceivable that the considerable machinery developed in the rational-choice framework might be brought to bear.

The empirical counterpart to the rational choice analysis of group decisions is to be found within the organizational psychology literature. Here, the dominant approach is to model such groups as information processing mechanisms. Therefore, this literature focuses on measuring the extent to which previously unshared information becomes shared, and on the basis of what sorts of information the ultimate decision is made [10-12]. The assumption is that a better decision will be made as more privately-held information is revealed. It is found that shared information drives decision outcomes, suggesting that sources of unshared information (or unique expertise) may be ignored by unstructured groups [13]. An analysis by Kameda et al. attempts to identify members in a group who are "cognitively central" – such members are able to translate otherwise unshared information into terms that other group members are familiar with [14]. Kameda's work suggests that information-sharing depends as much on the recipient's ability to digest it as it does on the speaker's willingness to share it. Nevertheless, the organizational psychology literature does not explicitly include institutional background of group members as an explanatory variable. Indeed, most experiments are performed using undergraduate student volunteers. In a real-world expert committee, cognitively-central members might be those committee members who are able to translate ideas across disciplinary boundaries. Nevertheless, the extension of this work to expert committees is not straightforward, precisely because the nature of expertise is that, while it is highly relevant to the problem at hand, it is not widely shared.

The approach taken within this paper is inspired by work within the anthropology and Science, Technology and Society (STS) literatures. In particular, the penetrating analyses of Mary Douglas note that social group membership affects perception of data [4]. Institutional membership is conferred upon those who structure categories of causality in a manner that is consistent with institutional norms [15]. This is reflective of a wider principle in anthropology that different cultures will selectively direct individuals' attention to the elements that are salient within their institutional structures. Among technical experts, this is reflected in the fact that each specialty possesses its own unique language and jargon, which carries with it an implicit scheme for categorizing perceived phenomena. On the other hand, an outsider to the institution, who is unfamiliar with the jargon used, may be unable to understand the discourse. This is because the specific jargon refers to commonly held sensory and social experiences that a member of another institution is unlikely to have directly encountered.

The STS literature extends this notion by noting that language is used as a cognitive mechanism to delineate professional boundaries. Within a specialty, word choice reflects the direction of an expert's attention toward a given interpretation of a problem that is consistent with that expert's training, while simultaneously directing that attention away from other possible interpretations [16-19]. The same institutions that drive selective perception and word choice also confer a sense of identity. March notes the dialectic between decision-making as rational choice based on a consequentialist pursuit of preferences and identity-based rule-following, ultimately noting that each viewpoint supplements the other [20]. We may therefore expect preferences to be correlated with institutional membership and, by extension, its associated jargon. This motivates an analysis of language in order to be able to examine institutional factors in group negotiation. This analysis is best performed on a relatively large corpus of technical committee meeting transcripts.

3. FDA Medical Device Approval

As a committee evaluation of a complex engineered system involving different specialties, the U.S. Food and Drug Administration (FDA) advisory panel meetings provide a rich data source from which we may study technical decision making by committees of experts [21]. Decisions made by technical expert committees in the FDA are analogous to those that must be made by committees of technical experts within a complex engineered system. As explained above, different experts possess different world-views, potentially leading to different, and equally legitimate, readings of uncertain evidence. Reaching a design decision requires that information from these different specialties be aggregated in some way. Ideally, the ultimate decision would be well-informed by all perspectives in the room.

3.A. Data Availability

One of the primary advantages to using the FDA Advisory Panels as a case study is the availability of data. There are 20 different panels whose transcripts are recorded over a period of ten years. This leads to the possibility of examining hundreds of committee meetings – a sufficiently large number that potentially generalizable findings may be inferred. If the study were to expand to include the drug-approval committees within the FDA, the number of cases upon which we may draw would number in the thousands.

The empirical analysis mentioned above requires data in the form of committee meeting transcripts. These are often not recorded in textual form, or are proprietary to the organization that commissioned the committee. We therefore turn to transcripts of expert committees that are a matter of public record. The ideal data source must have the following attributes:

1. Analysis or evaluation of a technological artifact
2. Participation of multiple experts from different fields or areas of specialization
3. A set of expressed preferences per meeting(such as a voting record)
4. Multiple meetings, so as to enable statistical significance

These requirements are met by the Food and Drug Administration's medical device advisory panels.

3.B. Multi-Stakeholder Environment

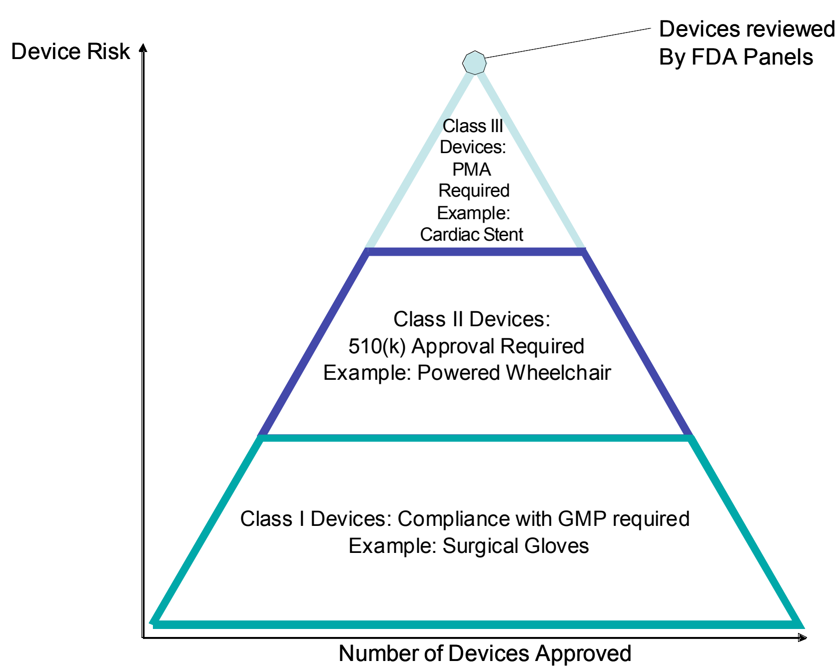

Figure 1: Medical devices are classified into three categories based upon risk to the patient. PMA = Pre-Market Approval; GMP = Good Manufacturing Practices.

The task of approving medical devices for the US market falls to the Food and Drug Administration's Center for Devices and Radiological Health (CDRH). Figure 1 provides an overview of the process by which a device is reviewed for approval by CDRH.

FDA divides medical devices into three classes, corresponding to patient risk. A subset of the most risky (i.e., Class III) devices, are reviewed by FDA advisory panel committees "as needed" [22]. Devices brought to committees for review are generally those that the FDA does not have the "in-house expertise" to evaluate. As such, the devices under evaluation by the committees are likely to be the most controversial innovations facing medical practice, and those facing the most uncertainty. Furthermore, advisory panel members are "by definition, the world's experts who are engaged in cutting-edge bench science, clinical research and independent consulting work" [21]. Advisory panels therefore serve to bring needed expert knowledge and political credibility to the FDA device approval process. Audience members routinely include representatives of the media, consumer advocate groups, the financial community, and competitor companies, all of whom seek information regarding how the medical device might perform on the market [23]. Panel recommendations, and the judgments and statements of individual members carry significant weight both within and outside the FDA.

3.C. Collaborative Technical Decision-Making in the FDA

As in a complex engineered system, technical experts in the FDA may not have an explicitly political aim. Nevertheless, their decisions may be perceived as biased by those who would have made a different decision in their place. Although FDA advisory committees are aimed at producing "evidence-based" recommendations, different interpretations of the evidence allow room for debate. Panel members' professional experiences might allow for intuition that can seem to go against the indications shown by the data. Friedman ([24]) expressed a concern that this constitutes a form of "specialty bias," especially when multiple specialties are involved. On the other hand, this view presupposes that a reading of the data that is entirely uninformed by past experience is best, which seems to obviate the role of expertise in advisory panel decision making. Others argue that conflicts of interest should be mitigated in advisory panels. On the other hand, a prominent study recently found only a minor correlation between conflict of interest and voting patterns, with no actual effect on device approval [25]. A distinction must be drawn between decision-making that is based on evidence and decision-making that is driven by one "orthodox" reading of the evidence.

4. Computational Approach

How are we to know if the panel members possess concurrent readings of the evidence? The work cited above suggests that the determination of institutional and other interpersonal affinity might be identified through the use of common language and jargon. If so, we might conclude that they find the same elements of the problem salient. This further suggests cultural affinity. The most direct way of operationalizing these insights is to attempt to cluster speakers by the co-occurrence patterns of their discourses. In particular, we use Bayesian modeling, to determine whether actors within a committee meeting are using similar terminology to discuss the common problem to be solved.

4.A. Previous Work

The choice of Bayesian topic models is driven by a desire to make the assumptions underlying this analysis minimal and transparent while maintaining an acceptable level of resolution of patterns inferred from the transcript data. An earlier iteration of this work, based on [26], used Latent Semantic Analysis (LSA) [27], a simpler predecessor of topic models, to study the same corpus of FDA transcripts. In practice, LSA can identify and separate major sources of variance in word choice within a discourse, as in evaluating the divergence between two groups of speakers (e.g., identifying device sponsors versus committee-members within the resulting latent semantic space). Nevertheless, LSA has some well-known limitations that stem from its use of Singular Value Decomposition (SVD) to analyze word co-occurrence patterns. Among these is the assumption that words are embedded within a Euclidean "semantic-space". This particular assumption breaks down when comparing words that are polysemous – i.e., having the same spelling but different meanings (compare "bat" the animal vs. "bat" in the context of baseball). LSA represents the location of these words in the Euclidean semantic space as the average over the two separate locations – an incorrect representation. Furthermore, LSA assumes that the noise around each word's location in the Euclidean space is normally-distributed, an assumption that introduces increasingly more distortion into the analysis as a given speaker uses fewer words. These limitations make it difficult to resolve the linguistic attributes of individual speakers, particularly in the absence of extensive speaker data within a given meeting. Furthermore, the latent dimensions of the LSA feature space, which nominally correspond to latent concepts of a discourse, are often difficult to interpret. These limitations motivate the use of a Bayesian model (for an excellent comparison of LSA to Bayesian models of text analysis, see [28]).

4.B. The Author-Topic Model

Given our assumption that each speaker possesses an institutional signature in his or her word choice, we would like to have the identity of the speaker inform the selection of topics. We therefore use a variant of Rosen-Zvi et al.'s Author-Topic (AT) Model [29], which creates probabilistic pressure to assign each author to a specific topic. Shared topics are therefore more likely to represent common jargon. Like LSA, AT also uses a term-document matrix as input. The AT model is implemented as follows:

4.B.1. Construction of a word-document matrix

We begin by parsing a committee-meeting transcript into a word-document matrix. Consider a corpus of documents, D, containing n documents $d_1 \ldots d_n$. Consider, as well, the union of all words over all documents:

$$W = \bigcup_{i=1}^{n} d_n \qquad (1)$$

Suppose there are m>n words in W, $w_1 \ldots w_m$. We may therefore construct a "word-document matrix", **Y**, with dimensions m x n, where each element in the matrix, y_{jk}, consists of a frequency count of the number of times word j appears in document k. For the analyses reported in this paper, a word-document matrix was constructed using the Python 2.5 programming language. Non-content-bearing "function words", such as "is", "a", "the", etc., were pre-identified and removed automatically. In addition, words were reduced to their roots using PyStemmer, a Python implementation of Porter's Snowball algorithm [30]. The resulting corpus generally consists of ~25,000 word tokens, representing about m = 2500 unique words in about n = 1200 utterances.

4.B.2. AT Model Structure and Implementation

Whereas LSA performs singular value decomposition on **X**, the Author-Topic model provides a more structured analysis. The Author-Topic model provides an analysis that is guided by the authorship data of the documents in the corpus, in addition to the word co-occurrence data used by LSA. Each author (in this case, a speaker in the discourse) is modeled as a multinomial distribution over a fixed number of topics that is pre-set by the modeler. Each topic is, in turn modeled as a distribution over words. A plate-notation representation of the generative process underlying the Author-Topic may be found in [29].

The Author-Topic model is populated using a Markov-Chain Monte Carlo Algorithm that is designed to converge to the distribution of words over topics and authors that best matches the data. Information about individuals authors is included in the Bayesian inference mechanism, such that each word is assigned to a topic in proportion to the number of words by that author already in that topic, and in proportion to the number of times that specific word appears in that topic. Thus, if two authors use the same word in two different senses, AT will account for this polysemy. Details of the MCMC algorithm implementation are given in [31]. The AT model was implemented in MATLAB using the Topic Modeling Toolbox algorithm [31].

4.B.3. Model Parameters

The AT model requires the selection of two parameters. Each author's topic distribution is modelled as having been drawn from a uniform Dirichlet distribution, with parameter α. A Dirichlet distribution is used because it is the conjugate prior of the multinomial distribution. One may think of α as a smoothing parameter. Values of α that are smaller than unity will tend to more closely fit the author-specific topic distribution to observed data. If α is too small, one runs the risk of overfitting. Similarly, values of α greater than unity tend to bring author-specific topic distributions closer to uniformity. A value of α=50/(# topics) was used for the results presented in this paper. This is the values suggested by the creators of the Topic Modelling Toolbox after extensive empirical testing. Similar to α is the second Dirichlet parameter, β, from which the topic-specific word distributions are drawn. β values that are large tend to induce very broad topics with much overlap, whereas smaller values of β induce topics which are specific to small numbers of words. Following the empirical guidelines set forth by Griffiths and Steyvers [31], and empirical testing performed by the authors, we set the value of β = 200/(# words).

4.B.4. AT Model Output

When applied to a transcript, we treat each utterance as a document. Thus, the meeting transcript may be viewed as a corpus. Words within each utterance are grouped into topics with probability proportional to the number of times that word has been previously used in that topic, and the number of times that word's "author" (i.e., speaker) has previously used that topic. Figure 2 shows sample output of the Author-Topic model applied to the FDA Meeting held on March 4^{th}, 2002.

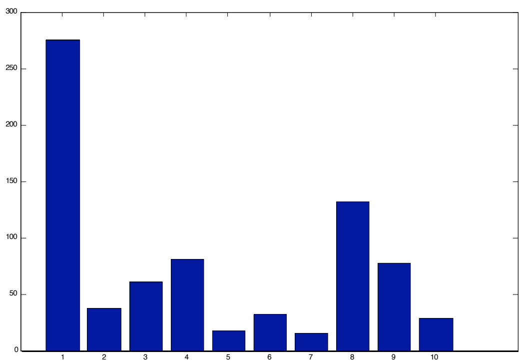

Figure 2: Sample output from the Author-Topic model run on the FDA Circulatory Systems Devices Advisory Panel Meeting for March 4th, 2002. This chart is the per-speaker topic distribution for one of the panel members

Table 1: The top five word-stems for one run of the AT model on the corpus for the Circulatory Systems Devices Panel Meeting of March 4, 2002

Topic Number	Top Five Word-Stems
1	'clinic endpoint efficaci comment base'
2	'trial insync icd studi was'
3	'was were sponsor just question'
4	'patient heart group were failur'
5	'devic panel pleas approv recommend'
6	'think would patient question don'
7	'dr condit vote data panel'
8	'effect just trial look would'
9	'lead implant complic ventricular event'
10	'patient pace lead were devic'

Each topic in Figure 2 may be identified by its most probable words. Table 1 displays the top five most probable word stems for each topic.

Within a clinical trial administered by the FDA, a device manufacturer must meet a certain set of clinical “endpoints”, often manifested as a proportion of a population that is free from disease or adverse events (e.g., device failure). Such trials typically have different endpoints for device safety and efficacy, both of which must be met. From this table, we can see that this panel member’s major focus involved questions of what was the appropriate clinical endpoint for the study in question (often a common debate on these panel meetings). This may indicate an interest in the efficacy endpoints (as opposed to the safety endpoints).

Speakers who speak often, and focus on one aspect of discourse will be more likely to have their words assigned to a topic focused on that speaker. If they focus on several aspects of the discourse in concert with other speakers (e.g., if they engage in a discussion), they will tend to have their words assigned to a number of topics related to their areas of focus. If they do not speak often, but are focused in their area of discourse, their words will likely be assigned to topics defined by other speakers. Finally, if they speak rarely, and are unfocused, their words will be assigned uniformly at random to all topics.

In practice, application of the AT model tends to group all committee voting members into the same topic. This occurs because the intra-committee variance in word usage is low compared to the word usage between the committee and the device sponsors, FDA representatives, etc. The AT model provides a convenient solution to this problem through the creation of a “false author” that is assigned to all committee members’ documents [32]. Thus, all words that are common to committee members are assigned to this false author, whereas those words that are unique to each member are preserved in their final topic distribution.

4.C. Generation of Social Networks

The above methodology can give us insight into the topics of interest for each speaker. Nevertheless, topics, on their own, provide little direct information about how individual speakers might relate to one another. Instead, we would like to use the topic information provided by the AT model to generate a social network.

4.C.1. Network Construction

We would like to link together speakers who commonly use the same topics of discourse. In particular, we examine each author-pair's joint probability of speaking about the same topic.

$$P(X_1 \cap X_2) = \frac{\sum_{i}^{T} P(Z = z_i \mid X_1) * P(Z = z_i \mid X_2)}{\sum_{i}^{T} \sum_{j}^{T} P(Z = z_i \mid X_1) * P(Z = z_j \mid X_2)} \qquad \textbf{(2)}$$

We would like to be able to construct an Author-Author matrix, **A**, with entries equal to 1 for each linked author pair, and entries equal to 0 otherwise. This may be interpreted as a social network [33].

4.C.2. Author-Author Matrix Determination

The AT model outputs an Author-Topic matrix that gives the total number of words assigned to each topic for each author. This information must be reduced to the **A** matrix identified above. Given no prior information about a given author's topic distribution, we might assume that that such a distribution is uniform over all topics. Therefore, we might expect *a priori* that the joint probability that any author pair would be linked would be uniform. In other words, if there are 10 topics, we would expect every author-pair to have a 10% probability of being linked, *a priori.* We consider an author pair to be linked within a given model iteration if that pair's joint probability exceeds what we would expect under a uniform distribution.

Each social network generated using this scheme is the result of one MCMC iteration. Multiple iterations, when taken together, form a probability distribution over a set of possible Author-Topic assignments, and therefore, feasible network topologies. We can expect that different iterations of the MCMC algorithm will yield different graphs.

Averaging over multiple MCMC iterations enables a social network with weighted links to be inferred, where the weight of each link is proportional to its frequency of occurrence among iterations. Nevertheless, the variability among draws from the MCMC algorithm suggests that links should not be weighted. Histograms of the distribution of these link frequency values tend to show a bimodal structure; suggesting that a description of author pairs as either connected or not connected is parsimonious.

In order to generate the networks shown in this paper, all author-pairs that were linked in more than 95% of all iterations were considered strongly-linked. All author-pairs that were linked in more than 90% of all iterations were considered weakly-linked.

5. Preliminary Network Analyses

The previous section demonstrated how social networks can be built. The following section begins a preliminary analysis of the capabilities of the methodology outlined in this paper.

5.A. Grouping by Medical Specialty

The results of the above analysis methodology support the assertion that language and medical specialty are correlated. Nevertheless, some meetings display voting along institutional lines more clearly than do others. For example, Figure 3 shows a strong grouping by medical specialty.

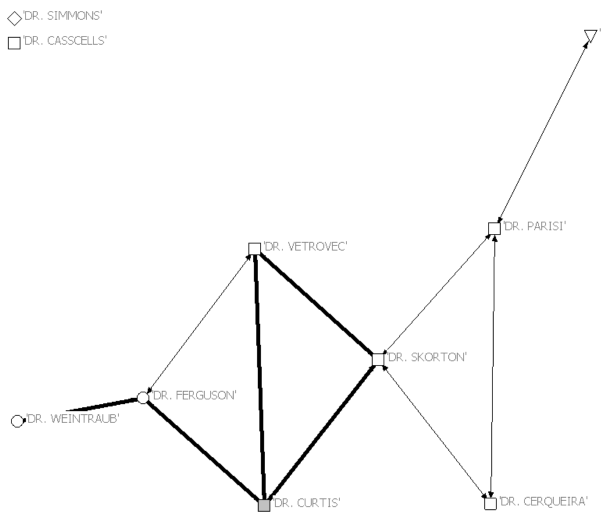

Figure 3: Graph of the FDA Circulatory Systems Advisory Panel meeting held on April 24, 1998. This meeting yielded a consensus approval of the medical device under analysis. Node shape represents medical specialty (circles are surgeons, squares are cardiologists, diamonds are electrophysiologists). Dr. Curtis, in grey, was the committee chair.

5.B. Grouping by Votes

In situations where the panel's vote is split, the method described in this paper can often isolate voting cliques (see Figure 4). In Figure 4, voting member cohesion dominates specialty cohesion. Nevertheless, this graph shows members of the same voting coalition to be connected. This suggests that the device reviewed in these meetings might be evoking an identity that transcends medical specialty.

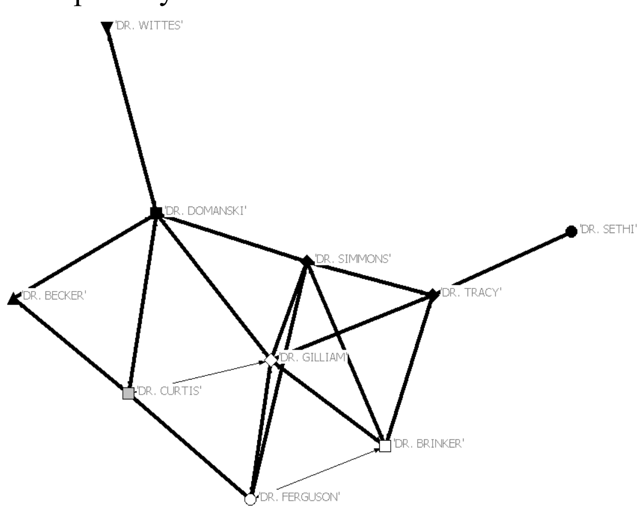

Figure 4: Graph of the FDA Circulatory Systems Advisory Panel meeting held on June 29, 1998. This device was not approved. Node shape represents medical specialty (circles are surgeons, squares are cardiologists, diamonds are electrophysiologists, down triangles are statisticians, and up triangles are neurologists). Dr. Curtis, in grey, was the committee chair. Non-approval votes are in black; approval votes are in white. In this meeting, vote is not correlated with medical specialty.

5.C. Idiosyncratic Choice

On June 23, 2005 the Circulatory Systems Devices Panel held a meeting to determine whether a particular device should be approved for a Humanitarian Device Exemption. Such a meeting likely appeals to a sense of personal ethical responsibility that transcends medical specialty. In situations such as these, we might expect that individual votes and connectivity patterns will be more idiosyncratic and exhibit less coherence. Figure 5 shows the connectivity pattern for this meeting. Note that this graph cannot be as easily partitioned by vote or by medical specialty, although yes voters do tend to congregate around the committee chair, whereas no-voters tend to congregate with the abstaining voters in the cluster near the upper-left. Why did these particular sets of experts choose to group together? Perhaps they shared a particular way of looking at the data that cut across medical specialty. Nevertheless, the fact that they did not all vote the same way indicates that there is an idiosyncratic element to their decision. Two people may agree on how to frame a problem, but may ultimately disagree on the correct course of action given that frame. This may be the case in those situations where approval voters are linked to large groups of non-approval voters and vice-versa.

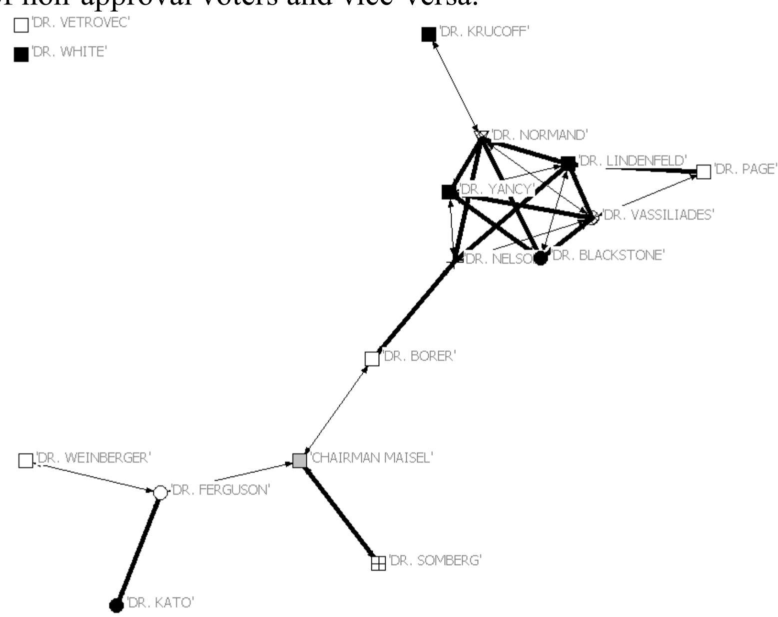

Figure 5: Graph of the FDA Circulatory Systems Advisory Panel meeting held on June 23, 2005. Node color represents the vote (black is against humanitarian device exemption, white is in favor of humanitarian device exemption, cross-hatched is abstention. Dr. Maisel, the committee Chairman, is grey). Node shape represents medical specialty (circles are surgeons, squares are cardiologists, crosshairs are a pharmacologist, down triangle is a statistician, plus is a bioethicist).

6. Conclusions

This work provides a preliminary computational methodology for the quantitative analysis of decisions by committees of technical experts. Ultimately, this research is intended to contribute along three lines:

1. **Methodological**: An algorithm and methodology for the analysis of expert committee decision making via the analysis of meeting transcripts.
2. **Theoretical**: New insights into group decision-making and socio-technical interactions.
3. **Practical**: Policy recommendations for how to structure approval committees to enable medical device safety and efficacy while still promoting innovation.

The method outlined in this paper relies on meeting transcripts to generate empirical findings regarding committee decision-making. It is seemingly sensitive to the limitation that not all committee members might express their views truthfully. Nevertheless, it is very difficult for individuals to avoid using jargon that they are familiar with. This is because word choice reflects a form of socio-cultural identity, as shaped by adherence to a professional institution.

Perhaps a larger concern is the inability of the method to differentiate between agreement and argument. Two actors are linked if they discuss the same topics. They may do so because they agree on some aspect of the device review, or because they are debating over interpretation of a given element of the debate. This is evident in the figures shown above, where linkage does not always indicate voting similarity. The committee chair, whose job it is to foster consensus, is likely going to attempt to frame the problem in such a way that the frame and the outcome is uniformly agreed-upon. If agreement on outcomes is not easily obtainable, fostering an ambiguous frame or multiple frames may help to reach consensus on voting outcomes (cf. [34]). Future work will focus on further elaborating this framework and on the determination of valence on each of the links in the graph. Determining signs for these graph links will enable a more direct comparison of the voting record to the graphical structure.

Acknowledgements

The author would like to thank Prof. Chris Magee, Prof. Maria Yang, and Dr. Joseph Coughlin for their extensive discussions and valuable mentorship, and the MIT-Portugal Program for its generous support.

References

[1] D.A. Broniatowski, "A Method for Generating Social Networks from Meeting Transcripts," *International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Modeling Intercultural Collaboration and Negotiation (MICON)*, Pasadena, CA: Springer-Verlag, 2009.

[2] D.A. Broniatowski, M. Yang, J.F. Coughlin, and C.L. Magee, "Bayesian Analysis of Decision Making in Technical Expert Committees," Conference on Systems Engineering Research, Loughborough, UK: 2009, p. 10.

[3] D.A. Broniatowski, C.L. Magee, J.F. Coughlin, and M. Yang, "Quantitative Analysis of Group Decision Making for Complex Engineered Systems," IEEE Systems Conference, Vancouver, British Columbia, Canada: 2009.

[4] M. Douglas, How Institutions Think, Syracuse, New York: Syracuse University Press, 1986.

[5] C.D. Elder and R.W. Cobb, The Political Uses of Symbols, New York: Longman, 1983.

[6] P.R. Carlile and C.B. Schoonhoven, "A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development," Organization Science, vol. 13, 2002, pp. 442-455.

[7] P. Boffey, "Scientists and Bureaucrats: A Clash of Cultures on FDA Advisory Panel," Science, vol. 191, 1976, pp. 1244-1246.

[8] W. Gaertner, A Primer in Social Choice Theory, Oxford University Press, 2009.

[9] B. Visser and O.H. Swank, "On Committees of Experts," Quarterly Journal of Economics, 2007, pp. 337-372.

[10] S. Mohammed and E. Ringseis, "Cognitive Diversity and Consensus in Group Decision Making: The Role of Inputs, Processes, and Outcomes," Organizational Behavior and Human Decision Processes, vol. 85, Jul. 2001, pp. 335, 310.

[11] R. Klimoski and S. Mohammed, "Team Mental Model: Construct or Metaphor?," Journal of Management, vol.

20, Apr. 1994, pp. 403-437.
[12] S. Mohammed, “Toward an Understanding of Cognitive Consensus in a Group Decision-Making Context,” Journal of Applied Behavioral Science, vol. 37, Dec. 2001, pp. 408-425.
[13] J.R. Winquist and J.R. Larson, “Information Pooling: When It Impacts Group Decision Making,” JOURNAL OF PERSONALITY AND SOCIAL PSYCHOLOGY, vol. 74, 1998, pp. 371-377.
[14] T. Kameda, Y. Ohtsubo, and M. Takezawa, “Centrality in Sociocognitive Networks and Social Influence: An Illustration in a Group Decision-Making Context,” Journal of Personality and Social Psychology, vol. 73, 1997, pp. 309, 296.
[15] J.A. Brown, “Professional language: words that succeed,” Radical History Review, vol. 34, 1986, pp. 33-51.
[16] M. Mulkay, T. Pinch, and M. Ashmore, “Colonizing the Mind: Dilemmas in the Application of Social Science,” Social Studies of Science, vol. 17, 1987, p. 231.
[17] R. Rapp, “Extra chromosomes and blue tulips: medico-familial interpretations,” Living and Working with the New Medical Technologies: Intersections of Inquiry, 2000.
[18] L. Winner, “On not hitting the tar-baby,” The Whale and the Reactor: A Search for Limits in an Age of High Technology. Chicago, vol. 111, 1986, pp. 138-154.
[19] C. Cohn, “Sex and Death in the Rational World of Defense Intellectuals,” Signs, vol. 12, 1987, pp. 718, 687.
[20] J.G. March, A Primer on Decision Making: How Decisions Happen, New York, NY: The Free Press, 1994.
[21] L.A. Sherman, “Looking Through a Window of the Food and Drug Administration: FDA's Advisory Committee System,” Preclinica, vol. 2, 2004, pp. 99-102.
[22] S. Parisian, FDA Inside and Out, Front Royal, VA: Fast Horse Press, 2001.
[23] W.L. Pines, How to Work With the FDA: Tips from the Experts, Washington, DC: FDLI, 2002.
[24] R.S. Friedman, “Representation in Regulatory Decision Making: Scientific, Industrial, and Consumer Inputs to the F.D.A.,” Public Administration Review, vol. 38, 1978, pp. 205-214.
[25] P. Lurie, C.M. Almeida, N. Stine, A.R. Stine, and S.M. Wolfe, “Financial Conflict of Interest Disclosure and Voting Patterns at Food and Drug Administration Drug Advisory Committee Meetings,” Journal of the American Medical Association, vol. 295, 2006, pp. 1921-1928.
[26] D.A. Broniatowski, C.L. Magee, J.F. Coughlin, and M. Yang, “The Influence of Institutional Background on the Approval of Engineered Systems,” Redondo Beach, CA: 2008, pp. 1-11.
[27] S. Deerwester, S.T. Dumais, G.W. Furnas, T.K. Landauer, and R. Harshman, “Indexing by Latent Semantic Analysis,” Journal of the American Society for Information Science, vol. 41, 1990, pp. 391-407.
[28] T.L. Griffiths, M. Steyvers, and J.B. Tenenbaum, “Topics in Semantic Representation,” PSYCHOLOGICAL REVIEW-NEW YORK-, vol. 114, 2007, p. 211.
[29] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, “The author-topic model for authors and documents,” Proceedings of the 20th conference on Uncertainty in artificial intelligence, Banff, Canada: AUAI Press, 2004, pp. 487-494.
[30] M.F. Porter, Snowball: A language for stemming algorithms. October 2001.
[31] T.L. Griffiths and M. Steyvers, “Finding scientific topics,” Proceedings of the National Academy of Sciences of the United States of America, vol. 101, Apr. 2004, pp. 5228-5235.
[32] M. Rosen-Zvi, T. Griffiths, P. Smyth, and M. Steyvers, “Learning Author Topic Models from Text Corpora,” Nov. 2005.
[33] S. Wasserman and K. Faust, Social Network Analysis: Methods and Applications, Cambridge University Press, 1994.
[34] R.W. Cobb and C.D. Elder, Participation in American Politics: The Dynamics of Agenda-Building, Baltimore and London: The Johns Hopkins University Press, 1983.