Quantitative Analysis of Group Decision Making for Complex Engineered Systems

David Andre Broniatowski, Joseph Coughlin

Engineering Systems Division
Massachusetts Institute of Technology
Cambridge, MA 02141
david@mit.edu

Christopher L. Magee, Maria Yang

Engineering Systems and Department of Mechanical
Engineering
Massachusetts Institute of Technology
Cambridge, MA 02141

Abstract—**Understanding group decision-making processes is crucial for design or operation of a complex system. Unfortunately, there are few experimental tools that might contribute to the development of a theory of group decision-making by committees of technical experts. This research aims to fills this gap by providing tools based on computational linguistics algorithms that can analyze transcripts of multi-stakeholder decision-making entities. The U.S. Food and Drug Administration medical device approval committee panel meetings are used as a data source. Preliminary results show that unsupervised linguistic analyses can be used to produce a formal network representation of stakeholder interactions.**

Keywords- multi-stakeholder decision-making, computational analysis

I. INTRODUCTION

The design of a sufficiently complex engineered system is beyond the capacity of a single human mind to understand. This is especially true for complex integrated systems and systems-of-systems (SoS), which Maier distinguishes by their stakeholder diversity, and requisite need for collaboration [1]. In general, any engineering system requires expert specialists who can comprehend the inner workings of specific systems and their components. As complexity grows, the cognitive capacity of a single system architect becomes insufficient, requiring that domain experts and system architects be able to pool their knowledge and learn from one another. 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 [2]. The mechanisms by which this training is achieved include acculturation within specific professional societies. Necessary to this process is the learning of that professional institution's language and jargon [3]. This leads to a situation wherein individual experts develop different views of the system. Without successful integration of these disparate views, important interactions between system components may not be captured. In such cases, the system becomes a boundary object, knowledge about which must be properly constructed by the experts in question, whereas institutions that support multi-stakeholder decision-making become integrating devices [4]. If such integration devices are to be effective, they must take into account the social dynamics that such decision-making bodies are likely to exhibit.

This problem is exacerbated by the fact that complex technical systems are often commissioned by multiple policy actors. Many of these actors have their own agendas and may selectively interpret, or "deconstruct and reconstruct" technical results to advance their individual policy goals [5]. The existence of different sources of expertise and institutional backgrounds among different types of engineers enables policy-makers to select technical biases that will support their final goals. This reflects a fundamental disconnect between the decision-making paradigms of the policy and engineering communities [6]. One could conceive of a scenario wherein policy can be made that is incongruent with the technical capabilities of the system, such as in the case of the International Space Station – whose original policy requirements were mutually technically inconsistent from scientific experimental, and orbital dynamics perspectives [7].

Understanding the social dynamics that underlie multi-stakeholder decision-making for complex engineered systems is therefore a crucial aspect of the design process. 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 [8]. 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. For example, successful committees routinely employ technical experts from multiple, potentially conflicting, organizations in order to resolve interoperability problems. 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., [9] 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 they interact with the technical specifics of the system in question.

This line of research is aimed at tracing how the topics of interest to individual committee members are integrated into a group decision. Application of the results of these analyses can aid in the design and development of a concordance between social structure and technical architecture that is aimed at

ensuring the continued value delivery of the technical system in question despite its complexity.

II. LITERATURE REVIEW

Much work has been focused on understanding how diverse teams of technical experts combine information [10-16]. The construct of a "mental model" has been used to describe how each expert views the object under analysis, with the overall understanding that the convergence of mental models may be necessary for an informed consensus decision to reached [15]. This research has made important contributions to our understanding of group decision-making dynamics. Nevertheless, there is currently no standard basis by which the different theories of group decision-making might be compared so as to determine which theories fit which empirical situations. This lack of a unifying theoretical base makes comparison across experiments and generalizations beyond laboratory settings difficult, since it is not clear which experimental phenomena map to which real-world situations. The analysis of decision-making by committees of technical experts in real-world scenarios has not been studied. Within the domain of economics, some studies have focused on generating theories of committee decision-making based on rational-choice [8, 17]. These theories make the simplifying assumption that all experts are the same. Furthermore, these theories have not been empirically validated. This research aims to fill a hole in the literature by providing a quantitative, repeatable methodology for the analysis of committees of technical experts in real-world settings.

Linguistic cues can provide insight into a particular expert's interests in a committee setting. For example, an expert's choice to define a problem a certain way (a "safety issue" vs. a "lack of confidence in the data") implicitly suggests a particular solution concept. Furthermore, group membership may affect perception of data, and each institution or specialty possesses its own unique language and jargon [2]. This is particularly true in technical disciplines, where conceptual precision is required to communicate within the specialty. Nelson notes the importance of written and oral language as a means of encapsulating and transferring tacit knowledge [18]. On the other hand, an outsider to the institution may be unable to understand the implications of that particular framing of the problem. Casting "organization [as] the mobilization of bias", Cobb and Elder recognizes institution-specific symbolism in language, noting that the choice of terminology in defining a problem may be seen as a means of mobilizing support [19]. Choosing to use specialized technical words serves to narrow the range of subjective meaning of otherwise ambiguous terminology (such as "safety" or "efficacy" in FDA's context) thereby implicitly redefining the problem according to a given speaker's particular interest and domain of expertise. Determining the speaker's intention in using both precise and "symbolic" language can allow insight into their preferences on the committee, as well as which institutional background might be driving their decision. This research aims to take advantage of these linguistic regularities in its construction of an analysis methodology.

III. METHODOLOGICAL APPROACH

Linguistic analysis based on a sufficiently large number of transcripts to generate robust results requires an automated approach. The first steps in this direction have been taken by Dong, who studied engineering teams' convergence on a shared design through a computational analysis of transcripts of their verbal exchanges [20]. Dong's computational approach, based on the Latent Semantic Analysis (LSA) algorithm [21], uses a metric semantic coherence as a proxy for team agreement [22-23]. Latent Semantic Analysis performs singular value decomposition on a matrix of word counts within each utterance in a meeting. The resulting singular vectors are rough approximations of the "concepts" of which each utterance is constructed. More details about LSA may be found in Broniatowski et al., which extended this approach to the analysis of expert panels in the Food and Drug Administration [24].

The LSA approach is limited by the assumptions underlying its model of discourse [25]. Approaches based on Bayesian inference, such as Latent Dirichlet Allocation (LDA), avoid these problems by modeling each utterance as made up of a set of topics, each of which is a multinomial probability distribution over the full set of words within a particular meeting [26]. A variant of LDA, the Author-Topic (AT) model, can been used to generate a distribution over topics for each participant in a meeting [27]. Application of the AT model proceeds as follows:

A. Construction of a word-document matrix

Consider a corpus of documents, $\mathcal{D}$, containing n documents $d_1 \ldots d_n$. Consider, as well, the union of all words over all documents, $\mathcal{W}$. Suppose there are m<n words, $w_1 \ldots w_m$. We may therefore construct a "word-document matrix", $\mathbf{X}$, with dimensions m x n, where each element in the matrix, x_{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 the PyStemmer algorithm.

B. AT Model Structure and Implementation

Whereas LSA performs singular value decomposition on $\mathbf{X}$, the Author-Topic model provides a more structured analysis. In particular, each author (in this case, a speaker in the discourse) is modelled as a distribution over topics, where each topic is, in turn modelled as a distribution over words. A plate-notation representation of the generative process underlying the Author-Topic model is found in Figure 1. The Author-Topic model is populated by 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. Details of the MCMC algorithm implementation are given in [27]. The AT model was implemented in MATLAB using the Topic Modelling Toolbox algorithm [25].
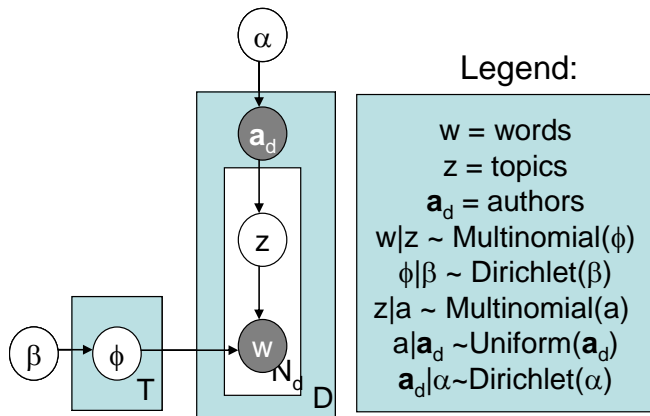
Figure 1. A plate notation representation of the Author-Topic model from [27]. Authors are represented by a multinomial distribution over topics, which are in turn represented by a multinomial distribution over all words in the corpus.

Legend:

w = words
z = topics
$\mathbf{a}_d$ = authors
$w|z \sim \text{Multinomial}(\phi)$
$\phi|\beta \sim \text{Dirichlet}(\beta)$
$z|a \sim \text{Multinomial}(a)$
$a|\mathbf{a}_d \sim \text{Uniform}(\mathbf{a}_d)$
$\mathbf{a}_d|\alpha \sim \text{Dirichlet}(\alpha)$

C. 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. 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. These different types of speakers are summarized in Table 1.

TABLE I.   DIFFERENT TYPES OF SPEAKERS IDENTIFIED BY THE AT MODEL. A FREQUENT, FOCUSED SPEAKER TENDS TO DRIVE TOPIC FORMATION, WHEREAS A RARE, UNFOCUSED SPEAKER TENDS TO BE ASSIGNED TO TOPICS DEFINED BY OTHERS. MULTI-FOCUS SPEAKERS MAY SERVE AS MEDIATORS.
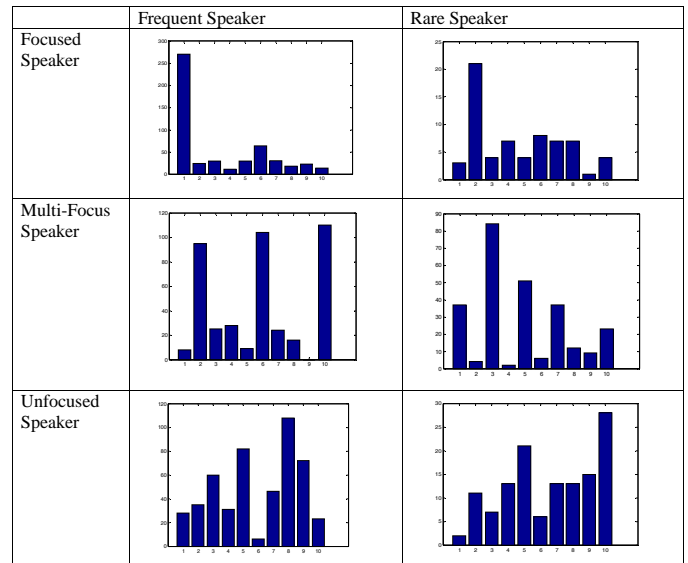
In practice, application of the AT model tends to group all committee voting members into the same topic. This occurs because the intra-committee variance in word usage is low compared to the word usage between the committee and the device sponsors, FDA representatives, etc. The AT model provides a convenient solution to this problem through the creation of a "false author" that is assigned to all committee members' documents. 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.

IV. CASE STUDY: FDA ADVISORY PANEL

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 [28].

A. Multi-Stakeholder Environment

As in a decision involving different stakeholders in a complex engineered system, the FDA decision-making process is embedded in a policy environment. 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, sourced from [29], provides an overview of the process by which a device is reviewed for approval by CDRH.
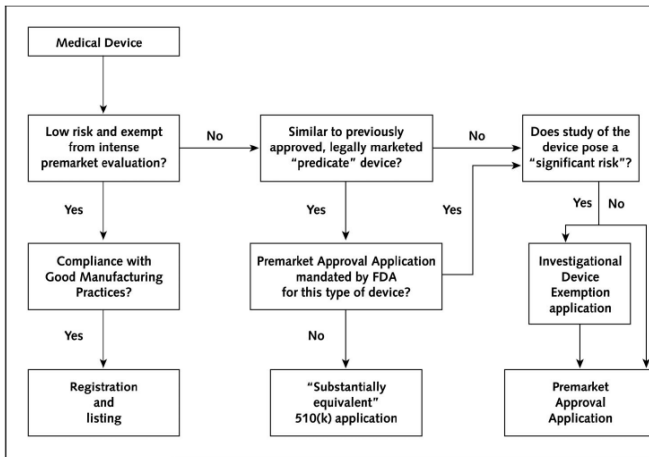
Figure 2.   Medical devices are classified into three categories based upon risk to the patient. Diagram sourced from [29]

The grant of a 510(k) or Pre-Market Approval (PMA) by the FDA allows a device to be marketed in the United States. These approvals often act as *de facto* monopolies for the device involved because any competitor must demonstrate additional safety or efficacy of the new device as compared to the initial baseline in order to receive approval. Advisory panels review devices "as needed" [30]. Devices brought to committees for review are generally those which 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 radical 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" [28]. Advisory panels therefore serve to bring needed expert knowledge and political credibility with industry and consumer advocate groups to the FDA device approval process. Audience members will include representatives of the media, consumer advocate groups, the financial community, and competitor companies, all of whom are looking for information regarding how the medical device might perform on the market [31]. Panel recommendations, and the judgments and statements of individual members, therefore carry significant weight both inside and outside the FDA.

B. Collaborative Technical Decision-Making in the FDA

As in a complex engineered system, technical experts in the FDA may not have a 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, differential interpretation of the evidence allows room for debate, and concomitant accusations of bias. Panel members' professional experiences might allow for intuition that can seem to go against the indications shown by the data. [32] 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 [33]. 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.

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 reading 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.

C. 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 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.

V.   PRELIMINARY APPLICATIONS AND RESULTS

A. Identification of Topics of Interest to Each Speaker

The AT model may be used to identify topics that are salient to each speaker. This can be helpful in determining how a coalition is built. Consider the meeting of the Circulatory Systems Devices Panel held on November 20, 2003. This meeting was convened to review a PMA for approval of the Taxus ® Paclitaxel Drug-Eluting Stent, designed and marketed by Boston Scientific Corporation. Taxus was the second drug-eluting stent on the market, following the panel's decision to approve Cordis Corporation's Cypher Sirolimus-Eluting Stent one year prior. The ultimate outcome of the meeting was a consensus decision to approve the PMA. The vast majority of decisions to approve a device come with conditions of approval that the panel recommends to the FDA that the sponsor must meet before the device can be marketed. This is in some ways analogous to requirements recommendations made by committees of designers of a SoS. The conditions of approval for the Taxus stent were as follows:

1. The labeling should specify that patients should receive an antiplatelet regimen of aspirin and clopidogrel or ticlopidine for 6 months following receipt of the stent.

2. The labeling should state that the interaction between the TAXUS stent and stents that elute other compounds has not been studied.

3. The labeling should state the maximum permissible inflation diameter for the TAXUS Express stent.

4. The numbers in the tables in the instructions for use that report on primary effectiveness endpoints should be corrected to reflect the appropriate denominators.

5. The labeling should include the comparator term "bare metal Express stent' in the indications.

Each of these conditions may be traced to a particular voting member, or set of voting members, on that panel, using the AT model. Table II, below, outlines the primary topics for each voting member. The top five words, identifying each voting member's preferred topic (out of 10 total), are identified, in addition to the topic proportion – the proportion of words spoken by that voting member that was assigned to that topic. Finally, each topic is assigned to a condition of approval as listed above.

TABLE II. RESULTS OF THE AUTHOR-TOPIC MODEL APPLIED TO A TRANSCRIPT OF THE CIRCULATORY SYSTEMS DEVICES PANEL MEETING OF NOV. 20, 2003. TOPICS CORRESPOND TO CONDITIONS OF APPROVAL FOR THE FINAL VOTE.

Committee Member	Major Topic of Interest (stemmed)	Topic Proportion	Corresponding Condition #
DR. HIRSH-FELD	'metal bare express restenosi paclitaxel'	0.36	5
DR. WHITE	'physician stainless ifu steel plavix'	0.42	1
DR. SOMBERG	'metal bare express restenosi paclitaxel' 'materi drug interact effect potenti'	0.30 0.29	5 2
DR. NORMAND	'tabl detail denomin six number'	0.56	4
DR. MORRISON	'metal bare express restenosi paclitaxel'	0.23	5
DR. YANCY	'drug clinic present appear event'	0.23	2
DR. WEINBERGER	'angiograph reduct nine think restenosi'	0.12	<None>
DR. MAISEL	'millimet length diamet coronari lesion'	0.34	3
DR. AZIZ	'know bit littl take present'	0.23	<None >

B. Generation of Social Networks

We may use the output of the Author-Topic model to gain insight into the social structure of a given committee. Since the results of the Author-Topic model assign each word to a topic, we may compare topic distributions across speakers. In particular, if two speakers' words are assigned to the same topic frequently (i.e., more than 20% of the time, for a model with ten topics), we say that they are "linked". Using authors as nodes, and the links derived from their topic distributions as edges, we may generate an author-topic graph. We refer to this graph as an *interpretation* of the meeting. One such interpretation is shown in Figure 3 below.
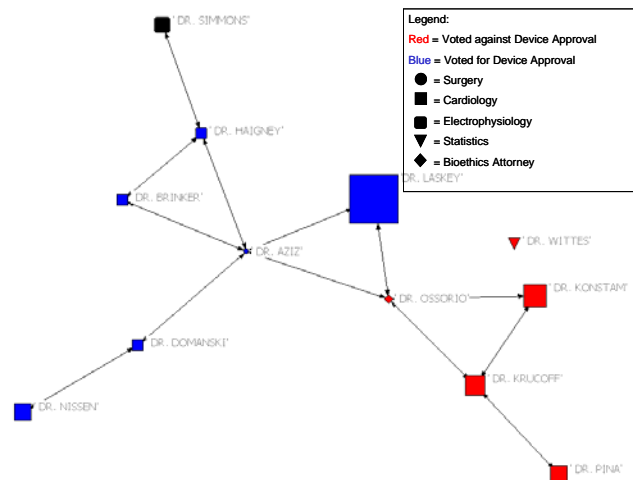

Figure 3. One interpretation of the meeting of the FDA Circulatory Systems Devices Advisory Panel Meeting held on March 5, 2002. Node size is proportional to the number of words spoken by the corresponding speaker. Dr. Simmons is a panel member who left the meeting before the vote occurred. Random seed = 3.14. Graphs were generated using UCINET.

Due to the probabilistic nature of the Author-Topic model's algorithm, we may expect different interpretations of a given meeting to occur based upon differing draws from the Author-Topic model's distribution over topic assignments for each word. A second interpretation for the same meeting as shown above is shown in Figure 4.
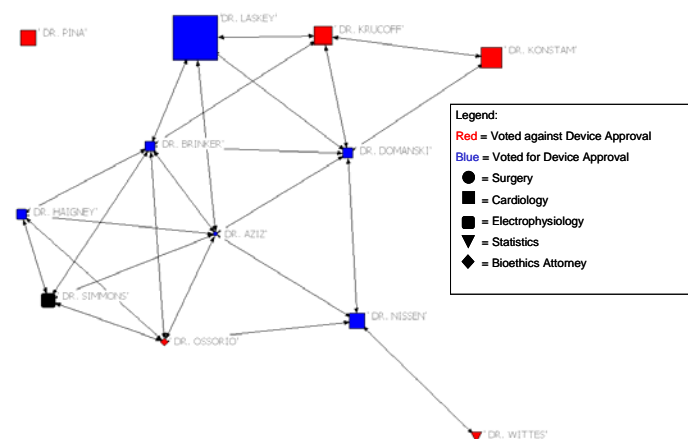

Figure 4. A second interpretation of the meeting of the FDA Circulatory Systems Devices Advisory Panel Meeting held on March 5, 2002. Node size

is proportional to the number of words spoken by the corresponding speaker. Random seed = 613. Graphs were generated using UCINET.

Each interpretation yields a connectivity pattern for each meeting. We may average these patters over multiple different interpretations, i.e., we average over the presence or absence of a link for each possible speaker pair. We may therefore get a measure for link strength. If a link occurs frequently across different interpretations, we say that it is a strong link. For example, a strong link might occur in more than half of the interpretations. A weak link is a link which occurs more often than average, but less frequently than a strong link. Finally, a spurious link is one which occurs less often than the average link. The result is displayed in Figure 5 below.
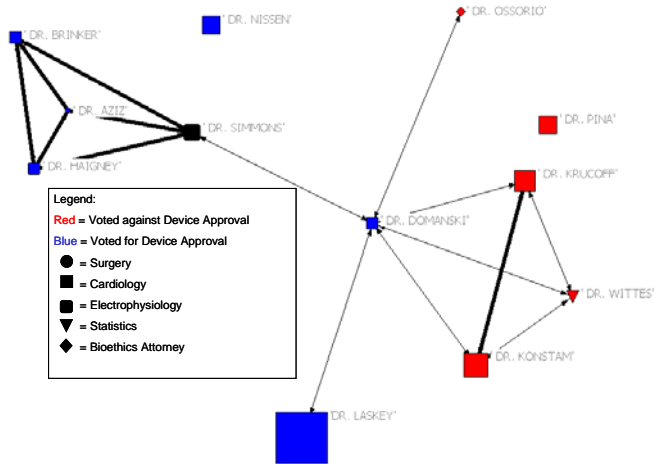

Figure 5.   Average of 200 interpretations for the meeting of the FDA Circulatory Systems Devices Advisory Panel Meeting held on March 5, 2002. A heavy line indicates a strong link (linked in >100 interpretations). A light line indicates that the speakers are linked more than the global average of all speakers. Spurious links have been eliminated

Drawing more samples from this distribution will necessarily yield a representation that is more statistically-robust. Based on the 200-sample draw shown above, we can already see the presence of "cliques" – i.e., connected sub-graphs – that correspond to like-minded voters. It appears that, in this particular case, committee-members generally voted sincerely. It is interesting that, in this particular case, every vote "counted" – i.e., if one committee member were to switch sides, the meeting's outcome would be changed. Such a situation has been postulated to support sincere voting over strategic voting [34].

C.   Next Steps

Future work will focus on analyzing these social networks and other derivatives of the analysis of transcripts using computational linguistics tools. For example, we might be able to identify speaker roles as follows: Speakers who are strongly linked to one another are likely to share a topic in common, and perhaps have similar jargon. Such a group of strongly-linked speakers may form a *coalition*, or may embody an argument. Speakers who are weakly linked to a large number of speakers

likely do not have very strong topic specificity, most likely due to infrequent speech. These speakers likely do not contribute much to the discussion and may vote more on the basis of exogenous factors. Speakers who are weakly linked to a small number of speakers, or not linked at all, yet who speak often, likely discuss a small number of topics that no one else shares. These speakers seem linguistically isolated, and therefore unlikely to have influenced others with their arguments. Finally, speakers who speak infrequently and are sparsely linked do not seem to have participated much in the formation of a group consensus.

VI.   CONCLUSIONS

The analysis of multi-stakeholder decision-making is a key element in the successful design and fielding of complex integrated systems. Nevertheless, we currently do not have any computational tools that can enable the creation and testing of a theory for how best to construct such committees. This research aims to fill this gap by providing a tool based on the computational analysis of meeting transcripts. Empirical insights from this analysis can yield valuable contributions to the theory of group decision-making, ultimately leading to decisions that better incorporate the perspectives of different participating technical experts while avoiding biased decision-making. This paper takes the first steps in that direction.

ACKNOWLEDGMENT

REFERENCES

[1]   Mark W. Maier, "Architecting principles for systems-of-systems," *Systems Engineering* 1, no. 4 (1998): 267-284, doi:10.1002/(SICI)1520-6858(1998)1:4<267::AID-SYS3>3.0.CO;2-D.

[2]   Mary Douglas, *How Institutions Think* (Syracuse, New York: Syracuse University Press, 1986).

[3]   Charles D. Elder and Roger W. Cobb, *The Political Uses of Symbols*, Longman professional studies in political communication and policy (New York: Longman, 1983).

[4]   P. R. Carlile and C. B. Schoonhoven, "A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development," *Organization Science* 13, no. 4 (2002): 442-455.

[5]   Sheila S. Jasanoff, "Contested Boundaries in Policy-Relevant Science," *Social Studies of Science* 17, no. 2 (1987): 195-230.

[6]   David André Broniatowski and Annalisa L Weigel, *Decision-Making in the Political and Technical Environments* (Cambridge, MA: Massachusetts Institute of Technology, 2006).

[7]   Howard E. McCurdy, The Space Station Decision : Incremental Politics and Technological Choice (The Johns Hopkins University Press , 1990).

[8]   Bauke Visser and Otto H. Swank, "On Committees of Experts," *Quarterly Journal of Economics* (2007): 337-372.

[9]   PM Boffey, "Scientists and Bureaucrats: A Clash of Cultures on FDA Advisory Panel," *Science* 191, no. 4233 (1976): 1244-1246.

[10]   Norbert L. Kerr, Robert J. MacCoun, and Geoffrey P. Kramer, "Bias in Judgment: Comparing Individuals and Groups," *Psychological Review* 103, no. 4 (1996): 687-719.

[11]   Tatsuya Kameda, Yohsuke Ohtsubo, and Masanori Takezawa, "Centrality in Sociocognitive Networks and Social Influence: An

Illustration in a Group Decision-Making Context," *Journal of Personality and Social Psychology* 73, no. 2 (1997): 309, 296.

[12] Susan Mohammed and Erika Ringseis, "Cognitive Diversity and Consensus in Group Decision Making: The Role of Inputs, Processes, and Outcomes," *Organizational Behavior and Human Decision Processes* 85, no. 2 (July 2001): 335, 310.

[13] K. J. Klein et al., "How do they get there? An examination of the antecedents of centrality in team networks," *Academy of Management Journal* 47, no. 6 (2004): 952-963.

[14] J. R. Winquist and J. R. Larson, "Information Pooling: When It Impacts Group Decision Making," *JOURNAL OF PERSONALITY AND SOCIAL PSYCHOLOGY* 74 (1998): 371-377.

[15] Richard Klimoski and Susan Mohammed, "Team Mental Model: Construct or Metaphor?," *Journal of Management* 20, no. 2 (April 1, 1994): 403-437, doi:10.1177/014920639402000206.

[16] Susan Mohammed, "Toward an Understanding of Cognitive Consensus in a Group Decision-Making Context," *Journal of Applied Behavioral Science* 37, no. 4 (December 1, 2001): 408-425, doi:10.1177/0021886301374002.

[17] R. K. Sah and J. E. Stiglitz, *Committees, Hierarchies and Polyarchies* (Department of Economics/Woodrow Wilson School of Public and International Affairs, 1987).

[18] Richard R. Nelson, "On the Uneven Evolution of Human Know-how," in *Technology, Institutions, and Economic Growth* (Cambridge, Massachusetts: Harvard University Press, 2005), 173-194.

[19] Roger W. Cobb and Charles D. Elder, *Participation in American Politics: The Dynamics of Agenda-Building* (Baltimore and London: The Johns Hopkins University Press, 1983).

[20] Andy Dong, Andrew W. Hill, and Alice M. Agogino, "A Document Analysis Method for Characterizing Design Team Performance," *Journal of Mechanical Design* 126 (2004): 378-385.

[21] Scott Deerwester et al., "Indexing by Latent Semantic Analysis," *Journal of the American Society for Information Science* 41, no. 6 (1990): 391-407.

[22] Andy Dong and A. Vande Moere, "Computationally assessing conceptual coherence in teamwork," Congress of the International Ergonomics Association (IEA'06) (2006).

[23] Peter W. Foltz, Walter Kintsch, and Thomas K. Landauer, "The Measurement of Textual Coherence with Latent Semantic Analysis," *Discourse Processes* 25, no. 2&3 (1998): 285-307.

[24] David André Broniatowski et al., "The Influence of Institutional Background on the Approval of Engineered Systems," in (presented at the CSER 2008, Redondo Beach, CA, 2008), 1-11.

[25] Thomas L. Griffiths and Mark Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences of the United States of America* 101, no. Suppl 1 (April 6, 2004): 5228-5235, doi:10.1073/pnas.0307752101.

[26] David M. Blei, Andrew Y. Ng, and Michael I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research* 3 (2003): 993-1022.

[27] Michal Rosen-Zvi et al., "The author-topic model for authors and documents," in *Proceedings of the 20th conference on Uncertainty in artificial intelligence* (Banff, Canada: AUAI Press, 2004), 487-494, http://portal.acm.org/citation.cfm?id=1036843.1036902.

[28] Linda Ann Sherman, "Looking Through a Window of the Food and Drug Administration: FDA's Advisory Committee System," *Preclinica* 2, no. 2 (2004): 99-102.

[29] William H. Maisel, "Medical Device Regulation: An Introduction for the Practicing Physician," *Annals of Internal Medicine* 140, no. 4 (2004): 296-302.

[30] Suzanne Parisian, *FDA Inside and Out* (Front Royal, VA: Fast Horse Press, 2001).

[31] Wayne L. Pines, *How to Work With the FDA: Tips from the Experts* (Washington, DC: FDLI, 2002).

[32] Robert S. Friedman, "Representation in Regulatory Decision Making: Scientific, Industrial, and Consumer Inputs to the F.D.A.," *Public Administration Review* 38, no. 3 (1978): 205-214.

[33] Peter Lurie et al., "Financial Conflict of Interest Disclosure and Voting Patterns at Food and Drug Administration Drug Advisory Committee Meetings," *Journal of the American Medical Association* 295, no. 16 (2006): 1921-1928.

[34] David Austen-Smith and Jeffrey S Banks, "Information aggregation, rationality, and the Condorcet Jury Theorem," *The American Political Science Review* 90, no. 1 (March 1996): 34.