Project Acronym:	**CATALYST**
Project Full Title:	**Collective Applied Intelligence and Analytics for Social Innovation**
Grant Agreement:	**6611188**
Project Duration:	**24 months (Oct. 2013 - Sept. 2015)**

D2.2 Analytics for Social Innovation Networks: Design Rationale

Deliverable Status:	**Draft**
File Name:	**CATALYST_ D2.2_WP2_UZH_V1.0_05022014.pdf**
Due Date:	**December 2013 (M3)**
Submission Date:	**February 2014 (M5)**
Dissemination Level:	**Public**
Task Leader:	**University of Zurich**

The CATALYST project consortium is composed of:

SO	Sigma Orionis	France
I4P	Imagination for People	France
OU	The Open University	United Kingdom
UZH	University of Zurich	Switzerland
EN	Euclid Network	United Kingdom
CSCP	Collaborating Centre on Sustainable Consumption and Production	Germany
Purpose	Purpose Europe	United Kingdom
Wikitalia	Wikitalia	Italy

Revision control

Version	Author	Date	Status
0.1	Mark Klein (University of Zurich)	January, 2014	Draft
0.2	Stéphanie Albiéro (Sigma Orionis)	February 04, 2014	Quality check
1.0	Stéphanie Albiéro (Sigma Orionis)	February 05, 2014	Peer review & Final version

Table of contents

Executive summary

The present document is a deliverable of the CATALYST project, funded by the European Commission's Directorate-General for Communications Networks, Content & Technology (DG CONNECT), under its 7th EU Framework Programme for Research and Technological Development (FP7).

This deliverable reports on Task 2.2, whose goal was to specify the analytics needed to achieve CATALYSTs mission of creating next-generation social innovation platforms. In the following the report will briefly review, as background, the weaknesses of current social innovation technologies as well as the approach that CATALYST is taking to address these problems. It then identifies the role that analytics play in CATALYST, and describes the methodology for identifying which analytics are needed. The analysis results themselves are available in the appendix.

1. Weaknesses with current social innovation technology

Humanity now finds itself faced with highly complex and often highly contentious challenges – such as climate change, the spread of disease, international security, scientific collaborations, product development, and so on - that call upon us to bring together large numbers of experts and stakeholders to innovate and deliberate collectively on how they can best be solved. Current social media technologies such as email, blogs, wikis, chat rooms, and web forums provide unprecedented opportunities for such interactions to take place, but have yet to realize their potential, running into serious challenges that include:

- *Platform Islands:* There are many social media platforms, splitting users into islands and thereby disrupting the free flow of ideas necessary for effective social innovation.
- *Cognitive Clutter:* Social media discussions produce large, redundant, and highly disorganized collections of contributions of widely varying quality, making it difficult to find the "good stuff" amongst all the noise.
- *Shallow Contributions:* Social innovation systems tend to generate large numbers of relatively shallow ideas, rather than a smaller number of deeply considered ones.
- *Unsystematic Coverage:* Current social innovation systems include no mechanism for ensuring that the ideas submitted comprehensively cover the different facets of the problem at hand, generally resulting in spotty coverage of the solution space.
- *Poor Idea Evaluation:* Social innovation systems do not currently provide effective techniques for helping the crowd identify the best ideas at large scales.
- *Poor visualization:* Existing social innovation systems provide only minimal tools for visualizing the key outputs of a deliberation, typically requiring that users simply read the whole corpus if they want a comprehensive picture of what has taken place.

2. Our approach: harvesting social media

The Catalyst project is about developing collective intelligence technologies that enable qualitatively more effective social innovation for complex and controversial problems, architected as follows:

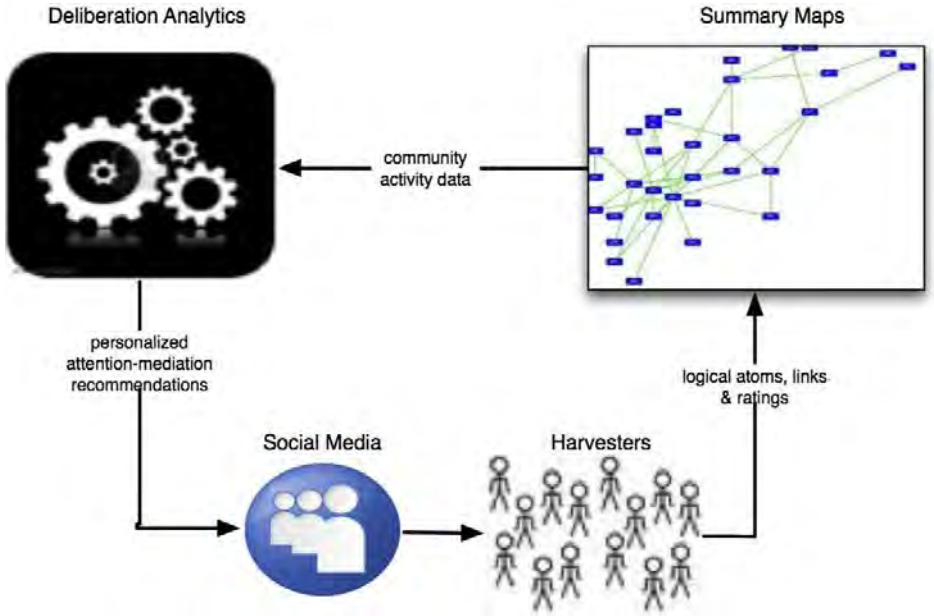

Figure 1. Architecture for the CATALYST social innovation system.

Vast online communities are already engaged in social innovation interactions using existing social media platforms such as email, web forums, blogs, Twitter, Facebook, and so on. Rather than attempting to supersede these platforms, our project *augments* them, using the core concepts of *harvesters*. "Harvesters" are communities of individuals who, supported by software tools, scan existing social media in order to harvest the most interesting and important issues, ideas, and arguments on a given topic and capture them in as organized, non-redundant summaries.

Summaries are represented as *argument maps*[1] (Buckingham Shum, 2003), which are tree structures made up of *issues* (questions to be answered), *ideas* (possible answers for a question), and pro/con *arguments* (statements that support or rebut an idea or argument):

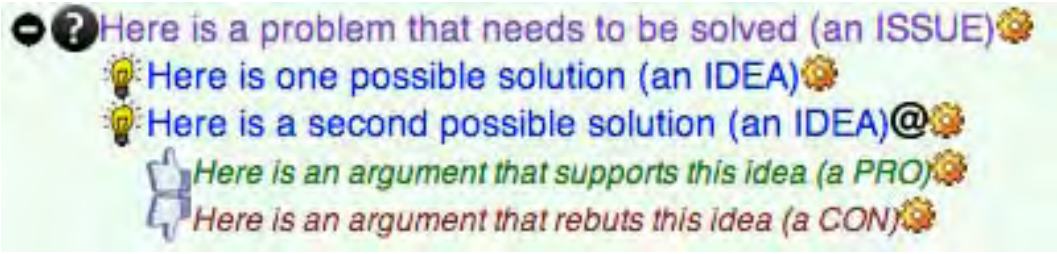

Figure 2. Example of an argument map.

The harvesting process provides a powerful approach for addressing the weaknesses of current social innovation technologies. Summary maps integrate contributions from multiple "platform islands", enabling the free flow of ideas needed for effective social innovation. The systematic structure of summary maps makes it much easier to see what has, and has not, been contributed so far, fighting cognitive clutter, enabling good visualization and systematic coverage, and facilitating collaborative refinement by making it clear how ideas can build upon one another.

[1] We will also, as the project proceeds, investigate other forms of 'knowledge cartography' (Okada *et al*, 2008) such as concept maps (Cañas & Novak, 2008), Delphi causality graphs (Linstone & Turoff 1975, Turoff *et al*. 1999; 2002), and richer ontologies (e.g. Buckingham Shum, et al., 2007; Rowe & Reed, 2008).

3. The Role of analytics

A critical challenge for making harvesting work is *attention allocation.* Even moderately complex societal challenges can involve scores of problems to solve, hundreds of possible solutions, and thousands of arguments for and against these possible solutions. How can the harvesters in our system known which particular topics are most in need of relevant material? How can managers of the harvesting process understand which areas are progressing well, and which may require some kind of intervention? How can the customers of the harvesting process know when a given part of a summary map is "mature" (i.e. comprehensively covers the key problems, solutions, and arguments) and thus ready to studied in detail?

The CATALYST project is meeting this challenge by developing *deliberation analytics*, i.e. algorithms that calculate deliberation metrics and map them to personalized attention mediation suggestions. If these algorithms work effectively, every social innovation participant can know *where their efforts can do the most good,* so the collective intelligence of the system is maximized.

4. Analytics identification methodology

The analytics under CATALYST are identified using *process-goal-exception* analysis, a technique developed by a member of the CATALYST team (Klein 2003). The key idea is that analytics can be viewed as the processes we put in place to identify, and respond, when a process deviates from its ideal functioning. This methodology allows us to identify process deviations and their associated responses in a systematic way that fosters complete coverage. It works as follows:

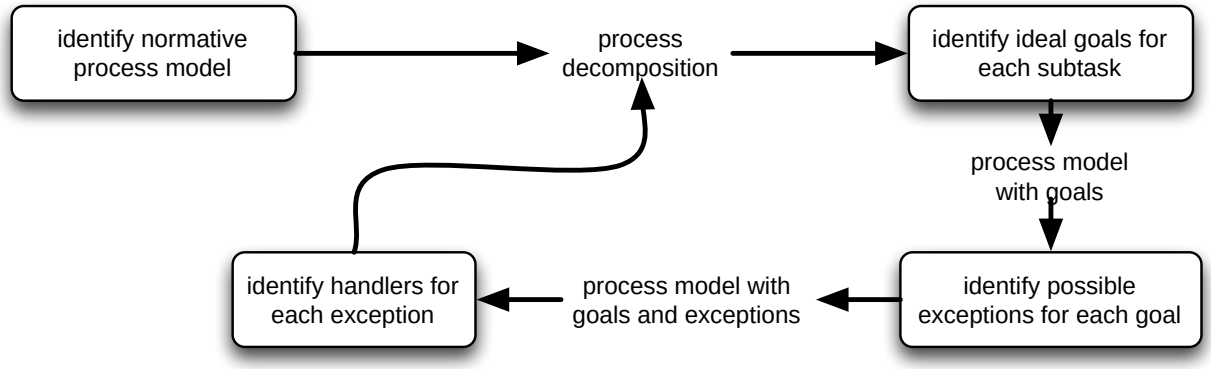

Figure 3. Process-Goal-Exception analysis, the methodology used for identifying analytics.

Identify normative process model: The first step is to identify a model of how the target process should work. The core process supported by the CATALYT system is social innovation. Our model of this process consists of the following subtasks (Walton and Krabbe, 1995) (Eemeren and Grootendorst, 2003):

Social Innovation Process					
1. Identify problems to solve	2. Identify possible solutions for these problems	3. Evaluate the candidate solutions	4. Select the best solution(s) from amongst the candidates	5. Enact the selected solution(s)	6. Learn from experience

This model is potentially iterative: enacting a selected solution (step 5) can, for example, lead the community to identify new problems to solve (step 1). Note also that social innovation engagements will necessarily include *all* these steps: it depends on the *purpose* of the engagement (Conklin, 2005), which can for example include:

- *Brainstorming*: create a list of solution options for a problem (step 2). Examples of this include strategic crowdsourcing in a company (before prioritization and decision by executive committee), or public consultation for a City. This can include using creativity techniques such as recombining known ideas
- *Argumentation*: debate the relative merit of competing solution options (step 3). This can include the use of simulation and forecasting tools to assess the probable impact of the options under consideration.
- *Decision -making*: select the preferred option from among a menu of alternatives (step 4).
- *Design enhancement*: refine an existing solution design (i.e. start with step 5, and then loop back to step 1).

The CATALYST social innovation process includes two key sub-processes. One is *harvesting*, wherein participants feed content, e.g. found in conventional social media, into the social innovation system. The harvesting sub-process consists of the following subtasks:

Harvesting Process			
Find interesting content	Summarize as map		
	unbundle	tag	organize

Harvesters find interesting content in social media platforms (such a Facebook, Twitter, mailing lists and blogs) where discussions about social innovations are taking place. This content is then parsed into "atoms" (i.e. individual issues, ideas, or arguments) tagged with their type (e.g. issue, idea, pro or con, evidence) and their topic area. These tagged atoms, in turn, are organized into summary maps.

The second key sub-process is *certification*, wherein moderators check the content contributed by authors in order to ensure it is organized so as to maximize it's ease-of-use for contributors and customers:

Certification Process							
Acquire post that needs attention	Check post for correctness				Take action on post		
	Bundling	Title	Location	Substance	[de-] certify	discard	repair

All new posts began with "pending" status, and become visible to the community at large only when certified. Moderators acquire posts that need attention (i.e. either pending posts, or certified posts that have been tagged as having problems) and then check the posts for correctness (i.e. whether they have substantive relevant content, are "unbundled" into individual issues ideas and arguments, have a clear title, and are placed in the correct part of the map). The moderator can then [de-]certify the post, discard it, and/or repair it.

Identify goals: The next step is to identify what each task in the process should ideally achieve: its' *goals*. Our current model of the social innovation process includes the following goals:

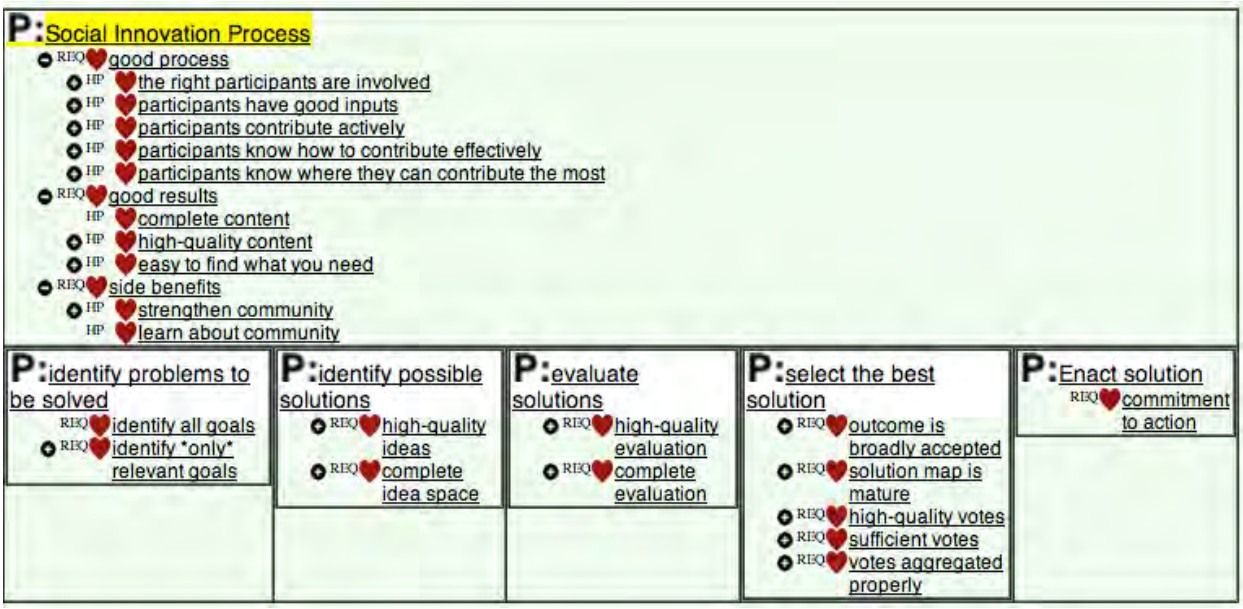

Figure 4: Top-level process model for social innovation.
P: *= process,* ♥ *= goal.*

A social innovation process should, for example, use a good process (i.e. where the right people contribute actively and effectively to performing the most critical tasks) to achieve good results (i.e. complete, high-quality, well-organized content) while also strengthening and learning about the members of the user community.

Identify exceptions: For each goal, we then identify how it can be violated (the *exceptions*). The goal of having the right participants involved, for example, can have the following exceptions:

Figure 5: Exceptions and handlers for a goal in the social innovation process

= exception = metric process = handler process

We can have too few authors, for example, or inadequate diversity in the author population.

Identify handlers: For each exception, finally, we identify handler processes that can (1) *detect* when the exception is taking place (i.e. via metrics), and (2) *resolve* that exception (i.e. via attention mediation interventions). We can *detect too few authors,* for example, using a metric that assess the width of the contribution activity histogram, and we *handle low author diversity* by encouraging participation from community members with underrepresented demographics (figure 5). Exception handler processes, like any other process, can themselves fail: the exception analysis process can be applied to handlers, just like any other process.

The attention-mediation interventions are personalized based on each participants' roles and past activity. The *customer* for a deliberation, for example, can be notified of topics that are mature and ready to be "harvested". A *topic manager* (responsible for ensuring a social ideation engagement achieves useful results) can be notified about which parts of the deliberation are dysfunctional (e.g. exhibit balkanization or groupthink). A *moderator* can be notified about users who consistently do (or do not) author well-structured and well-regarded posts, in order to inform training, moderator recruitment and/or rewards for top contributors. A *contributor* can be notified of content they can contribute to, such as pet ideas whose support has dropped, or posts where their ratings appear to exhibit an irrational bias. The *class* of the contributor (Preece and Shneiderman, 2009) (e.g. heavy contributors vs. peripheral users) should also likely impact which attention notification alerts they receive.

The results of process-goal-exception analysis are captured using the following structure:

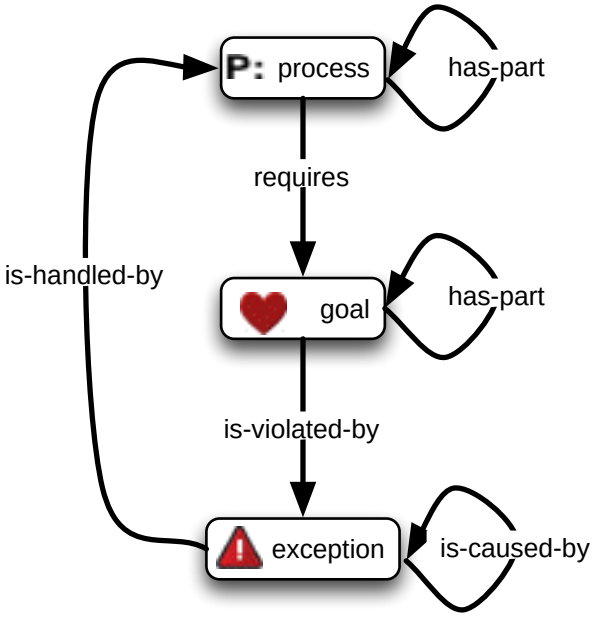

Figure 6. Entities and relationships that represent the results of process-goal-exception analysis.

Tasks in a process model are linked to their subtasks as well as to the goals they try to achieve. Goals are linked to their sub-goals, as well as to the exceptions that can violate them. Exceptions, finally, are linked to the other exceptions that may cause them, as well as to the (metric and attention mediation) processes that can detect and resolve these exceptions.

5. Analysis results

The CATALYST team has developed a substantive model of the social innovation process as well as of the analytics that can help make it work better. This model was created using a combination of top-down and bottom-up analysis:

- *Top-down*: The model incorporates insights from the research literature for such fields as organizational science, cognitive and social psychology, political and communication science, computational social science, computer-supported cooperative work, complexity science, and economics. The references section, below, lists many of the papers that were harvested for this analysis effort.
- *Bottom-up*: Task 2.1 collected extensive requirements, from the CATALYST community partners, on the limitations of existing social innovation systems ("pain points") and how they could be improved. These were also incorporated in our analysis.

Our team developed a web-based system to make it easy to view and update the results of the analytics identification process:

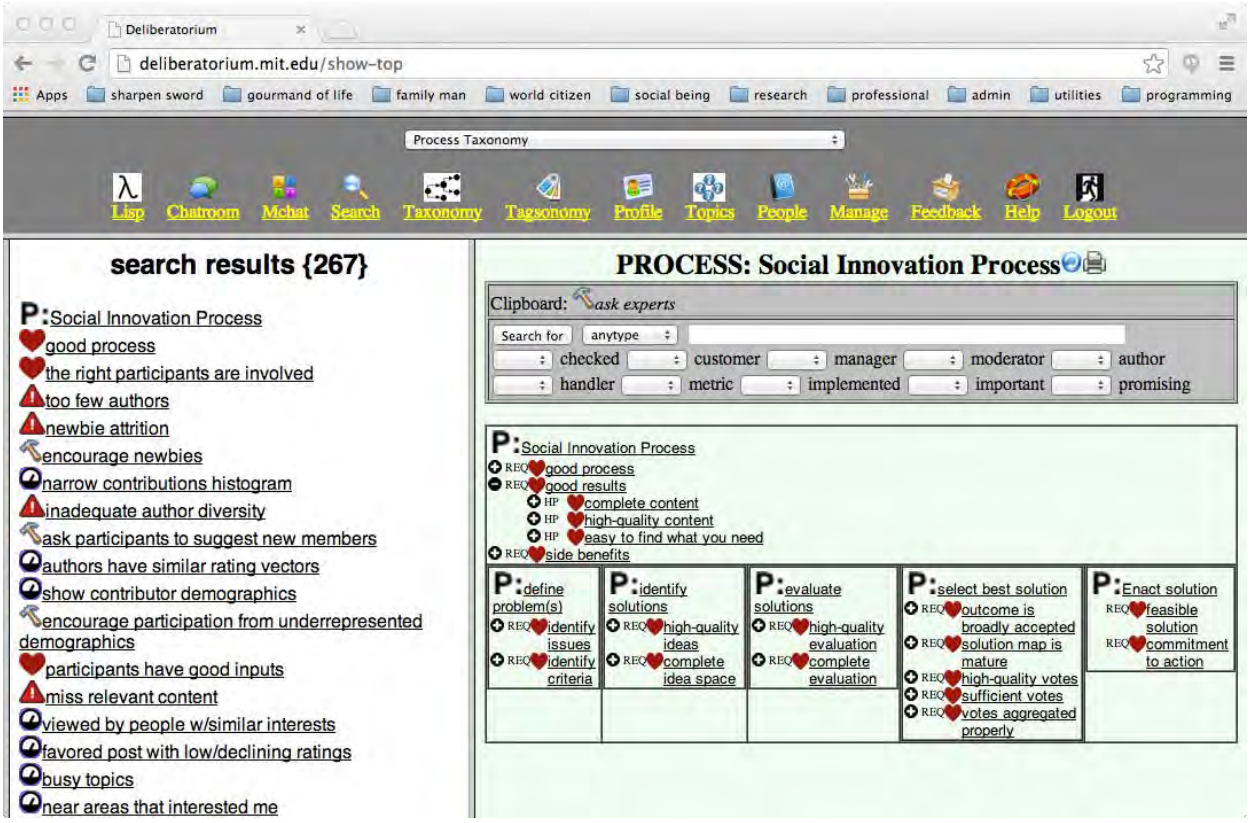

Figure 7. A screenshot of the web-based system for process-goal-exception analysis.

Users can click on the ⊕ and ⊖ icons to incrementally hide and reveal components of the analysis, and click on the components themselves to see more details on each one. A search capability allows users to find components with a given type and keywords.

Our analysis efforts have resulted, at the time of writing, in a model with nearly 300 components, with a particular focus on the first three steps of the social innovation process ("identify problems", "identify possible solutions", and "evaluate solutions") since these were identified as the most critical elements by the research and community partners. Every component in the model is tagged with the system role (author, moderator, manager, customer) it is relevant to, as well as with whether it was considered, at the time of writing, to be a promising candidate for early implementation in our system. The model is available in the Appendix to this document.

This deliberation metrics analysis should be viewed a *living document*. We will update it throughout the project as we develop an increasingly complete understanding of how to achieve more effective large-scale social innovation processes. We have developed a web-based collective intelligence system that we will use to gather feedback about our existing metrics, as well as suggestions for new metrics, from the argumentation-mapping community:

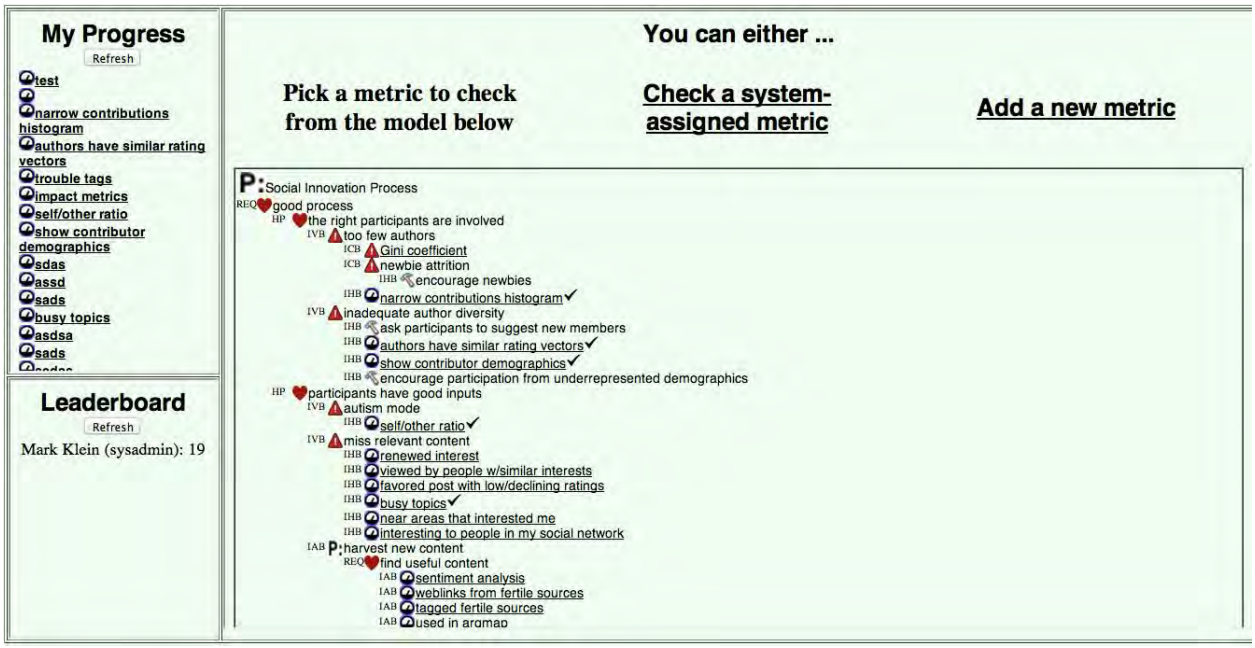

Figure 8. A screenshot of the web-based collective intelligence system for tapping the argument-mapping community's knowledge about deliberation metrics.

Potential communities for providing feedback on the metrics include the many users of such argument-mapping systems as Compendium, Cohere, Debategraph, Agora, Rationale, and ConsiderIt.

References

Accorroni, M., & Bentivoglio, C. A. (2009). *Supporting tutoring by improved statistical analysis of discussion forum*. 463-466.

Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. (2008). *Knowledge sharing and yahoo answers: everyone knows something*. Proceedings of the Proceeding of the 17th international conference on World Wide Web, 665-674.

Aldenderfer, M. S., & Blashfield, R. K. (1984). *Cluster Analysis.* Sage.

Almeida, R. B., Mozafari, B., & Cho, J. (2007). *On the evolution of wikipedia.*

Arackaparambil, C., & Yan, G. (2011). *Wiki-watchdog: Anomaly detection in Wikipedia through a distributional lens*. Proceedings of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Volume 01, 257-264.

Arazy, O., & Stroulia, E. (2009). *A utility for estimating the relative contributions of wiki authors.*

Austen-Smith, D., & Feddersen, T. J. (2006). Deliberation, preference uncertainty, and voting rules. *American Political Science Review*, *100(02)*(02), 209-217.

Baker, P. (2009). *Contemporary corpus linguistics.* Continuum Intl Pub Group.

Bakshy, E., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). *Everyone's an influencer: quantifying influence on twitter*. Proceedings of the ACM international conference on Web search and data mining, 65-74.

Baron, R. S., Kerr, N. L., & Miller, N. (1992). *Group process, group decision, group action.* Thomson Brooks/Cole Publishing Co.

Benevenuto, F., Magno, G., Rodrigues, T., & Almeida, V. (2010). *Detecting spammers on twitter.*

Benevenuto, F., Rodrigues, T., Almeida, V., Almeida, J., & Gonçalves, M. (2009). *Detecting spammers and content promoters in online video social networks*. 620-627.

Bian, J., Liu, Y., Zhou, D., Agichtein, E., & Zha, H. (2009). Learning to recognize reliable users and content in social media with coupled mutual reinforcement.

Biocca, F., Harms, C., & Gregg, J. (2001). *The networked minds measure of social presence: Pilot test of the factor structure and concurrent validity.*

Bolstad, W. M. (2010). *Understanding Computational Bayesian Statistics.* John Wiley.

Borgatti, S. P., Mehra, A., Brass, D. J., & Labianca, G. (2009). Network analysis in the social sciences. *science*, *323(5916)*(5916), 892.

Brandes, U., & Erlebach, T. (Eds.). (2005). *Network Analysis: Methodological Foundations.* Berlin, Heidelberg: Springer-Verlag.

Briggs, R. O., Reinig, B. A., Shepherd, M. M., Yen, J., & Nunameker, J. F. (1997). *Quality as a function of quantity in electronic brainstorming*. Proceedings of the System Sciences, 1997, Proceedings of the Thirtieth Hawaii International Conference on, 94-103 vol. 2.

Cappella, J. N., Price, V., & Nir, L. (2002). Argument Repertoire as a Reliable and Valid Measure of Opinion Quality: Electronic Dialogue During Campaign 2000. *Political Communication*, *19(1)*(1), 73 - 93.

Carrington, P. J., Scott, J., & Wasserman, S. (Eds.). (2005). *Models and Methods in Social Network Analysis.* Cambridge University Press.

Chai, K., Hayati, P., Potdar, V., Wu, C., & Talevski, A. (2010). Assessing post usage for measuring the quality of forum posts. *writing*, *3*, 10.

Conklin, J. (2005). *Dialogue Mapping: Building Shared Understanding of Wicked Problems.* John Wiley and Sons, Ltd.

Maciel, C., & Garcia, A. C. B. (2007). Design and Metrics of a 'Democratic Citizenship Community' in Support of Deliberative Decision-Making. *Electronic Government*, *4656*, 388-400.

Davies, T., & Chandler, R. (2011). Online Deliberation Design: Choices, Criteria, and Evidence. In T. Nabatchi, M. Weiksner, J. Gastil, & M. Leighninger (Eds.), *Democracy in Motion: Evaluating the Practice and Impact of Deliberative Civic Engagement.* Oxford University Press.

Diakopoulos, N., Naaman, M., & Kivran-Swaine, F. (2010). *Diamonds in the rough: Social media visual analytics for journalistic inquiry.*

Diakopoulos, N. A., & Shamma, D. A. (2010). *Characterizing debate performance via aggregated twitter sentiment*. Proceedings of the international conference on Human factors in computing systems, 1195-1198.

Forsyth, D. R. Group Dynamics Resource Page. https://facultystaff.richmond.edu/~dforsyth/gd/.

Du, N., Wu, B., Pei, X., Wang, B., & Xu, L. (2007). *Community detection in large-scale social networks*. Proceedings of the Web mining and social network analysis, 16-25.

Eemeren, F. H. v., & Grootendorst, R. (2003). *A Systematic Theory of Argumentation: The Pragma-dialectical Approach.* Cambridge University Press.

Ekstrand, M. D., & Riedl, J. T. (2009). *rv you're dumb: identifying discarded work in Wiki article history*. Proceedings of the

International Symposium on Wikis and Open Collaboration, 4.

Eysenbach, G. (2011). Can tweets predict citations? Metrics of social impact based on Twitter and correlation with traditional metrics of scientific impact. *Journal of Medical Internet Research*, *13(4)*(4).

Fiore, A. T., Tiernan, S. L., & Smith, M. A. (2002). *Observed behavior and perceived value of authors in usenet newsgroups: bridging the gap*. Proceedings of the Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves, 323-330.

Fisher, D., Smith, M., & Welser, H. T. (2006). *You are who you talk to: Detecting roles in usenet newsgroups*. Proceedings of the System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on, 59b-59b.

Fong, P. K. F., & Biuk-Aghai, R. P. (2011). *Visualizing author contribution statistics in Wikis using an edit significance metric*. Proceedings of the International Symposium on Wikis and Open Collaboration, 197-198.

Freelon, D. G. (2010). Analyzing online political discussion using three models of democratic communication. *new media & society*, *12(7)*(7), 1172.

Freelon, D., Watanabe, M., Busch, L., & Kawabata, A. (2008). *Town Halls of the Digital Age: Controversy and Ideology in Online Deliberation (and Beyond)*. Proceedings of the Conference of the Association of Internet Researchers, 350-353.

Ghosh, R., & Lerman, K. (2010). Predicting influential users in online social networks. *J arXiv preprint arXiv:1005.4882*.

Gonzalez-Bailon, S. (2010). The structure of political discussion networks: a model for the analysis of online deliberation. *Journal of Information Technology*, *25*, 230-243.

Graham, T., & Witschge, T. (2003). In search of online deliberation: Towards a new method for examining the quality of online discussions. *Communications*, *28(2)*(2), 173-204.

Havre, S., Hetzler, B., & Nowell, L. (2002). ThemeRiverTM: In search of trends, patterns, and relationships. *IEEE Transactions on Visualization and Computer Graphics*, *8(1)*(1), 9-20.

Heer, J., Mackinlay, J., Stolte, C., & Agrawala, M. (2008). Graphical histories for visualization: Supporting analysis, communication, and evaluation. *Visualization and Computer Graphics, IEEE Transactions on*, *14(6)*(6), 1189-1196.

Heymann, P., Koutrika, G., & Garcia-Molina, H. (2007). Fighting spam on social web sites: A survey of approaches and future challenges. *Internet Computing, IEEE*, *11(6)*(6), 36-45.

Ho, C. H., & Swan, K. (2007). Evaluating online conversation in an asynchronous learning environment: An application of Grice's cooperative principle. *The Internet and higher education*, *10(1)*(1), 3-14.

Hong, L., & Davison, B. D. (2010). *Empirical study of topic modeling in twitter*. 80-88.

Jain, A. K., Murty, M. N., & Flynn, P. J. (1999). Data clustering: a review. *ACM Computing Surveys*, *31(3)*(3).

Jamali, M., & Abolhassani, H. (2006). *Different aspects of social network analysis*. 66-72.

Java, A., Song, X., Finin, T., & Tseng, B. (2007). *Why we twitter: understanding microblogging usage and communities*. Proceedings of the WebKDD and SNA-KDD, 56-65.

Jones, R. (2011). Using metrics to describe the participative stances of members within discussion forums. *Medical Internet Research*, *13(1)*(1), e3.

Khurana, U., Nguyen, V. A., Cheng, H. C., Ahn, J., Chen, X., & Shneiderman, B. (2011). *Visual analysis of temporal trends in social networks using edge color coding and metric timelines*. Proceedings of the IEEE international conference on social computing (socialcom), 549-554.

King, I., Li, J., & Chan, K. T. (2009). *A brief survey of computational approaches in social computing*. Proceedings of the International Joint Conference on Neural Networks, 1625-1632.

Kittur, A., Suh, B., Pendleton, B. A., & Chi, E. H. (2007). He says, she says: conflict and coordination in Wikipedia. *SIGCHI Conference on Human Factors in Computing Systems*.

Knoke, D., & Yang, S. (2008). *Social network analysis* ((7)). Sage Publications, Inc.

Korfiatis, N., & Naeve, A. (2005). Evaluating wiki contributions using social networks: A case study on wikipedia. *Growth*.

Lange, K. (2010). Singular Value Decomposition. *Numerical Analysis for Statisticians,* 129-142.

Lee, K., Caverlee, J., & Webb, S. (2010). *Uncovering social spammers: social honeypots+ machine learning*. 435-442.

Lin, Y. R., Sundaram, H., De Choudhury, M., & Kelliher, A. (2009). *Temporal patterns in social media streams: Theme discovery and evolution using joint analysis of content and context*. Proceedings of the International Conference on Multimedia and Expo, 1456-1459.

Liu, B. (in press). Sentiment analysis: A multi-faceted problem. *IEEE Intelligent Systems*.

Lüdeling, A., & Kytö, M. (2008). *Corpus linguistics: An international handbook.* W. de Gruyter.

Klein, M. (2012). Enabling Large-Scale Deliberation Using Attention-Mediation Metrics. *Computer-Supported Collaborative Work*, *21(4)*(4), 449-473.

Mautner, G. (2009). Checks and balances: How corpus linguistics can contribute to CDA. *Methods of Critical Discourse Analysis. London: Sage,* 122-143.

Mehler, A. (2008). Large text networks as an object of corpus linguistic studies. *Corpus Linguistics: An International Handbook, Mouton de Gruyter, Berlin,* 328-382.

Mercier, H., & Landemore, H. (2012). Reasoning is for arguing: Understanding the successes and failures of deliberation. *Political Psychology*, *33(2)*(2), 243-258.

Muhlberger, P. (2006). Report to the Deliberative Democracy Consortium: Building a Deliberation Measurement Toolbox [HP Author1] Version 2; 12/06/2007.

Muhlberger, P., & Stromer-Galley, J. (2009). *Automated and hand-coded measurement of deliberative quality in online policy discussions*. Proceedings of the Proceedings of the 10th Annual International Conference on Digital Government Research: Social Networks: Making Connections between Citizens, Data and Government, 35–41.

Nelson, B. A., Wilson, J. O., Rosen, D., & Yen, J. (2009). Refined metrics for measuring ideation effectiveness. *Design Studies*, *30(6)*(6), 737-743.

Nisbet, D. (2004). Measuring the Quantity and Quality of Online Discussion Group Interaction. *Journal of eLiteracy*, *1*, 122-139.

Novak, T. P., & Hoffman, D. L. (1997). New metrics for new media: toward the development of Web measurement standards. *World Wide Web Journal*, *2(1)*(1), 213-246.

Pal, A., & Counts, S. (2011). *Identifying topical authorities in microblogs*. Proceedings of the international conference on Web search and data mining, 45-54.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, *2(1-2)*(1-2), 1-135.

Preece, J. (2001). Sociability and usability in online communities: determining and measuring success. *Behaviour & Information Technology*, *20(5)*(5), 347-356.

Preece, J., Nonnecke, B., & Andrews, D. (2004). The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, *20(2)*(2), 201-223.

Preece, J., & Shneiderman, B. (2009). The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on Human-Computer Interaction*, *1(1)*(1), 13-32.

Vries, R. D., Stanczyk, A., Wall, I. F., Uhlmann, R., Damschroder, L. J., & Kim, S. Y. (2010). Assessing the quality of democratic deliberation: A case study of public deliberation on the ethics of surrogate consent for research. *Social Science & Medicine*, *70 (12)*(12), 1896-1903.

Robert, C. P. (2001). *The Bayesian Choice – A Decision-Theoretic Motivation.*

Russell, M. (2011). *Mining the Social Web: Analyzing Data from Facebook, Twitter, LinkedIn, and Other Social Media Sites.* O'Reilly Media.

Samet, H. (2006). *Foundations of Multidimensional and Metric Data Structures.* Morgan Kaufman.

Shamma, D. A., Kennedy, L., & Churchill, E. F. (2009). *Tweet the debates: understanding community annotation of uncollected sources*. Proceedings of the SIGMM workshop on Social media, 3-10.

Shin, Y., Gupta, M., & Myers, S. (2011). *Prevalence and mitigation of forum spamming*. Proceedings of the INFOCOM, 2309-2317.

Shum, S. B., Liddo, A. D., Iandoli, L., & Quinto, I. (in press). A Debate Dashboard to Support the Adoption of Online Knowledge Mapping Tools. *VINE Journal of information and Knowledge Management Systems*.

Smith, M. A., Shneiderman, B., Milic-Frayling, N., Mendes Rodrigues, E., Barash, V., Dunne, C. et al. (2009). *Analyzing (social media) networks with NodeXL*. Proceedings of the international conference on Communities and technologies, 255-264.

Spatariu, A., Hartley, K., & Bendixen, L. D. (2004). Defining and Measuring Quality in Online Discussions. *The Journal of Interactive Online Learning*, *2(4)*(4).

Steenbergen, M. R., Bachtiger, A., Sporndli, M., & Steiner, J. (2003). Measuring political deliberation: a discourse quality index. *Comparative European Politics*, *1(1)*(1), 21-48.

Stromer-Galley, J. (2007). Measuring deliberation's content: A coding scheme. *Journal of Public Deliberation*, *3(1)*(1), 12.

Sunstein, C. R. (2002). The law of group polarization. *Journal of political philosophy*, *10(2)*(2), 175-195.

Sunstein, C. R. (2006). *Infotopia: How Many Minds Produce Knowledge.* Oxford University Press.

Sureka, A. (2011). Mining user comment activity for detecting forum spammers in YouTube. *J arXiv preprint arXiv:1103.5044*.

Svartvik, J. (2007). Corpus linguistics 25+ years on. *Language and Computers*, *62(1)*(1), 11-25.

Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, *37(2)*(2), 267-307.

Trénel, M. (2004). Measuring the quality of online deliberation. Coding scheme 2.4. *Social Science Research Center Berlin, Germany. Available at: www. wz-berlin. de/online-mediation/files/publications/quod_2_4. pdf*.

Trénel, M. (2006). *Facilitating deliberation online: What difference does it make?* Proceedings of the *2nd Annual Conference on Online Deliberation*, 20-22.

Turner, T. C., Smith, M. A., Fisher, D., & Welser, H. T. (2005). Picturing Usenet: Mapping Computer-Mediated Collective Action. *Journal of Computer-Mediated Communication*, *10(4)*(4), 00-00.

Ulicny, B., Kokar, M. M., & Matheus, C. J. (2010). Metrics For Monitoring A Social-Political Blogosphere: A Malaysian Case Study. *Internet Computing, IEEE*, *14(2)*(2), 34-44.

Van Alstyne, M., & Brynjolfsson, E. (2005). Global village or cyber-balkans? Modeling and measuring the integration of electronic communities. *Management Science*, *51(6)*(6), 851-868.

Varlamis, I., Eirinaki, M., & Louta, M. (2010). *A study on social network metrics and their application in trust networks*. 168-175.

Velásquez, J. D., & González, P. (2010). Expanding the possibilities of deliberation: the use of data mining for strengthening democracy with an application to education reform. *The Information Society*, *26(1)*(1), 1-16.

Viegas, F. B., Wattenberg, M., & Dave, K. (2004). Studying cooperation and conflict between authors with history flow visualizations. *SIGCHI conference on Human factors in computing systems*.

Walton, D. N., & Krabbe, E. C. W. (1995). *Commitment in dialogue: Basic concepts of interpersonal reasoning.* Albany, NY: State University of New York Press.

Wasserman, S., & Faust, K. (1994). *Social Network Analysis: Methods and Applications.* Cambridge University Press.

Welch, C. (2010). Measuring Balkanization in Wikipedia.

Yang, J., & Counts, S. (2010). Predicting the speed, scale, and range of information diffusion in twitter. *Proc. ICWSM*.

Ye, S., & Wu, S. (2010). Measuring message propagation and social influence on Twitter. com. *Social Informatics,* 216-231.

Yilmaz, L. (2009). *On the synergy of conflict and collective creativity in open innovation socio-technical ecologies*. Proceedings of the Conference on Computational Science and Engineering, 502 - 508.

List of tables and figures

Annex - Detailed metrics analysis

Social Innovation Process

♥ = GOAL P: = PROCESS ⚠ = EXCEPTION ◔ = METRIC 🔨 = HANDLER

P:Social Innovation Process
This is the exception analysis for the argument-map mediated social innovation process used in the CATALYST project.

1. HAS-PART **P:define problem(s)**
Define the problem(s) that the social innovation engagement is supposed to solve.

1.1. REQUIRES ♥**identify issues**
Describe the issues that need to be solved e.g. "what can we do to solve climate change"?

1.1.1. IS-VIOLATED-BY ⚠**missing key issues**

1.1.1.1. IS-HANDLED-BY 🔨**ask experts**
ask experts to pre-populate map with all key issues

1.1.1.2. IS-HANDLED-BY ◔**expert evaluation**
Experts assess whether or not map includes all key issues.

1.2. REQUIRES ♥**identify criteria**
identify the attributes of a good solution to the problem e.g. "limit average global temperature rise to no more than 2 degrees celsius "

1.2.1. HAS-PART ♥**identify *only* relevant criteria**
Identify only criteria that are relevant/important for this problem - no false positives.

1.2.1.1. IS-VIOLATED-BY ⚠**irrelevant criteria**

1.2.1.1.1. IS-HANDLED-BY ◔**low rating**
Criterion has a low rating score.

1.2.2. HAS-PART ♥**identify all criteria**
Identify all relevant/important criteria for this problem.

1.2.2.1. IS-VIOLATED-BY ⚠**parochial criteria**
The participants identify only criteria that impact their own self-interest and not common good criteria such as social welfare and equality.

1.2.3. HAS-PART ♥**identify *only* relevant criteria**
Identify only criteria that are relevant/important for this problem - no false positives.

1.2.3.1. IS-VIOLATED-BY ⚠**irrelevant criteria**

1.2.3.1.1. IS-HANDLED-BY ◔**low rating**
Criterion has a low rating score.

1.2.4. HAS-PART ♥**identify all criteria**
Identify all relevant/important criteria for this problem.

1.2.4.1. IS-VIOLATED-BY ⚠**parochial criteria**
The participants identify only criteria that impact their own self-interest and not common good criteria such as social welfare and equality.

2. HAS-PART **P:identify solutions**
Identify candidate solutions for the identified problems.

2.1. REQUIRES ♥**high-quality ideas**
Existing social media tend to elicit lots of shallow ideas, with highly variable quality and originality. How can we maximize the proportion of creative, high-quality, deeply considered ideas?

2.1.1. IS-VIOLATED-BY ⚠**idea sabotage**
People who don't like an idea edit it to make it worse.

2.1.1.1. IS-HANDLED-BY ◔**edit wars by partisans**
i.e. where someone who doesn't like idea is editing it in conflict with someone who likes the idea. In other words, look for alternating edits by users that appear to have divergent opinions (based on their rating behavior) about the issue they are proposing solutions for.

2.1.1.2. IS-HANDLED-BY ◔**hub-and-spoke interaction network**
Secretive sabotage communication patterns tend towards a hub-and-spoke architecture, as opposed to the network-topology connectivity that characterizes full open discussion. See: Brandy Aven (2011). The effect of corruption on organizational networks and individual behavior. Proceedings of the MIT WIDS colloquium (http//wids.lids.mit.edu/).

2.1.2. IS-VIOLATED-BY ⚠**solo ideation**
Authors do not collaborate to refine ideas.

2.1.2.1. IS-HANDLED-BY ◔**single editor**
This post has had only one editor (not counting moderators).

2.1.3. IS-VIOLATED-BY ⚠**insular ideation**
ideas do not build upon one another

2.1.3.1. IS-HANDLED-BY ◔**no common ground vocabulary**
look for growing use of shared words and word clusters within topics, which is a way of assessing whether people are building ideas by re-combining existing ones.

2.1.4. IS-ACHIEVED-BY ◔**low idea ratings**
The ideas receive low average "promising" rating from the community.

2.2. REQUIRES ♥**complete idea space**
We want to have a comprehensive picture of the most promising solutions for the problems focused on by the innovation engagement.

2.2.1. IS-VIOLATED-BY **incomplete idea coverage**
The deliberation has only incompletely covered the space of potentially relevant ideas for an (important) issue.

2.2.1.1. IS-HANDLED-BY **attention/importance ratio**
Measure the ratio of attention to issue importance for issues, and highlight the issues with particularly low scores. Issue importance can be calculated by accounting for the importance of the parent issues and promise of the parent ideas.

2.2.1.2. IS-HANDLED-BY **get n(idea) estimates**
Ask users to estimate how many good ideas there are for each issue (e.g. whenever someone creates an issue, or adds an idea to an issue, or even views an issue). The average of that number gives the standard, and we flag an exception if we are substantially below that number of ideas for an issue.

2.2.1.3. IS-HANDLED-BY **ask expert panel**
An expert panel assess whether or not the idea space is covered fully (for a given issue).

2.2.2. IS-VIOLATED-BY **creativity stagnation**
Few novel/interesting ideas are being generated as proposed solutions for a problem.

2.2.2.1. IS-CAUSED-BY **idea groupthink**
groupthink can be defined as a group dedicating the bulk of its attention to refining a single idea, often the first one endorsed by an influential figure, rather than comparing several alternatives in depth.

2.2.2.1.1. IS-HANDLED-BY **attention narrowing**
we can measure when one idea under an issue is receiving the bulk of the community's attention (views, rates, edits and additions) while competing ideas and their underlying arguments remain largely untouched

2.2.2.2. IS-HANDLED-BY **count # ideas rated as novel**
Participants assign a degree of novelty (High, Average and Low) to the posted ideas. The degree of novelty of an issue is the max value attributed to each idea associated to that issue. The degree of novelty of an idea is the average degree of novelty assigned by the crowd to the idea.

2.2.2.3. IS-HANDLED-BY **low vocabulary diversity**
We can measure the use of shared vocabulary in the ideas for a given issue. If there is heavy use of shared terminology, this suggests that the ideas are only moderately diverse. Ideas that are truly diverse will tend to use different vocabulary to express them. In other words, look for ideas that are quite different, in terms of the word frequency statistics, from the other ideas (for that part of the map). We can use LSA or LDA or other document similarity algorithms for this purpose.

2.2.2.4. IS-HANDLED-BY **idea GA**
Usually a solution consists of a *package* of interrelated ideas, so the complete solution space will consist of different combinations of "atomic" ideas. These recombinant space can of course be vast, however, so in practice we must focus on only "promising" packages if at all possible. We can use a GA approach to draw people towards posts. People can score posts on creativity vs. practicality, and weight creativity more at first, practicality more as we near the final point using latent semantic analysis to help identify out-of-box posts and give them higher fitness scores - to maintain diversity. The system can also point people to pairs of ideas - e.g. ideas for different parts of a system, or different ideas for the same subsystem – and suggest they create a new idea based on these existing ones. This has the advantage that we interleave generation and evaluation to help produce a more efficient process (as opposed to generate everything first, and then evaluate the whole redundant mess). The system can suggest users look at combinations that will speed that search for optimal idea combinations when issues are interdependent and utility functions are therefore nonlinear. This can be based on techniques for simulated annealing, creating sub-negotiations for tightly-interdependent issue clusters, etc.

2.2.2.5. IS-HANDLED-BY **"red herrings"**
Use "out of the box" prompts to help break a creative deadlock e.g.: • oblique strategy cards (phrases or cryptic remarks) • randomly selected ideas from the summary map • ideas selected from areas/people the author has heretofore ignored

3. HAS-PART **evaluate solutions**
evaluate solutions with respect to the goals identified for the deliberation

3.1. REQUIRES **high-quality evaluation**
The evaluation provides accurate assessments of the worth of proposed solutions.

3.1.1. HAS-PART **users understand content**
Users understand the content of the map well enough to offer informed evaluations of the ideas described therein.

3.1.1.1. IS-ACHIEVED-BY **narrative summaries**
Convert argmaps into easy-to-follow narrative summaries that make it easier for evaluators to see the key points they need to. Perhaps we can do so taking advantage of rhetorical structure theory?

3.1.2. HAS-PART **complete argumentation**
i.e. the evaluation includes a comprehensive overview of the arguments for and against each proposed solution idea

3.1.2.1. IS-VIOLATED-BY **missing arguments**
An idea is missing some important arguments for or against it.

3.1.2.1.1. IS-CAUSED-BY **self-focused**
The participants generate arguments that refer to criteria that concern them personally, but not those that impact welfare for other groups.

3.1.2.1.2. IS-HANDLED-BY **neglected criteria**
Few/no arguments have been created that evaluate a given idea with respect to one of the solution criteria for that problem.

3.1.2.1.3. IS-HANDLED-BY **few/no arguments**
i.e. there are (important) ideas that have few or no arguments attached to them

3.1.2.1.4. IS-HANDLED-BY **unbalanced arguments**
There is a large imbalance in the number of pros and cons in the debate over an idea.

3.1.2.1.5. IS-HANDLED-BY **few people contributed arguments**

3.1.2.1.6. IS-HANDLED-BY **idea/argument rating disconnect**
We can use such techniques as Bayesian inference (Bolstad, 2010) to propagate a user's ratings for arguments up the argument map to predict how the user should have rated the ideas these arguments address. If there is a large divergence between a user's predicted and actual ratings for an idea, that suggests that the user has not yet entered arguments that are compelling to him or her. A suggestion is to show users the possible misalignment between popularity ratings assigned freely by users and computed scores based on the structure of arguments and how much underlying pros and cons are supported. The gap can be used also as an incentive to users: a popular idea with poor computed score should invite its supporters to provide new arguments or improve the existing ones.

3.1.3. HAS-PART **high-quality argumentation**
The arguments entered in the summary map are well-founded.

3.1.3.1. IS-VIOLATED-BY **false premises**
the arguments made are based on false premises

3.1.3.1.1. IS-CAUSED-BY **argument sabotage**
Someone who disagrees with an argument edits it to sabotage it, rather than (properly) simply downrating or arguing against it.

3.1.3.1.1.1. IS-HANDLED-BY **argument edit wars**
assess the prevalence of edit wars (rapid alternating rollbacks) in the post edit histories – esp. by people who take differing positions on the issues they are editing arguments for. see ZIF workshop paper by János Kertész, Budapest: Edit wars on the Wikipedia:. an interesting measure for wikipedia article controversiality = # mutual reverts by more established user accounts (discounting young vandals)

3.1.3.1.2. IS-HANDLED-BY **low argument ratings**
An argument got a low average rating from the community.

3.1.3.1.3. IS-HANDLED-BY **expert evaluation**
Expert(s) judged that the argument is ill-founded.

3.1.3.2. IS-VIOLATED-BY **incorrect inference**
the arguments made are based on logical fallacies

3.1.3.2.1. IS-HANDLED-BY **automated feedback**
The use of artificial intelligence techniques holds promise to increase the effectiveness of argumentation systems by automatically analyzing user actions and providing supportive feedback. see: http//www.ascilite.org.au/ajet/ajet25/butchart.pdf Mclaren, B. M., Scheuer, O., & Mikšátko, J. (2010). Supporting collaborative learning and e-Discussions using artificial intelligence techniques. International Journal of Artificial Intelligence in Education, 20(1)(1), 1-46. Scheuer, O., McLaren, B. M., Loll, F., & Pinkwart, N. (2012). Automated Analysis and Feedback Techniques to Support Argumentation: A Survey. In: N. Pinkwart, & B. M. McLaren (Eds.), Educational Technologies for Teaching Argumentation Skills (pp. 71–124). Bentham Science Publishers.

3.1.3.2.2. IS-HANDLED-BY **expert judgment**
Expert(s) judge that the inference is faulty.

3.1.4. HAS-PART **high-quality ratings**
The community ratings for the argument map contents (i.e. whether issues are relevant, criteria are important, ideas are promising, arguments are compelling) are accurate.

3.1.4.1. IS-VIOLATED-BY **too few ratings**
There are too few ratings for a post to draw reliable conclusions about how the community judges it's value.

3.1.4.2. IS-VIOLATED-BY **dishonest ratings**
Ratings are dishonest.

3.1.4.2.1. IS-HANDLED-BY **rating inconsistency**
A participant gives inconsistent ratings e.g. they rate the arguments supporting an idea highly, but give the idea itself a poor rating.

3.1.4.3. IS-VIOLATED-BY **incorrect ratings**
the user's ratings for the posts are incorrect

3.1.4.3.1. IS-HANDLED-BY **missing suppoprt**
someone gave a strong + or - rating without a backup argument that they authored or highly rated. This can take several levels, according to the Discourse Quality Index (Jurgen Steiner): (1) The speaker does not present any arguments (asks, for example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

3.1.4.3.2. IS-HANDLED-BY **ignored arguments**
User did not attend to (read, rate) arguments when rating posts impacted by these arguments

3.1.4.3.3. IS-HANDLED-BY **irrational ratings**
see how well a model of rational rating predicts user's ratings. We can use knowledge about map structure to infer what a user's rating for a post "should" be given their ratings for posts below it in tree. For example, with mutex ideas, a rate for X is probably a rate against alternative to X.

3.1.5. HAS-PART **users understand content**
Users understand the content of the map well enough to offer informed evaluations of the ideas described therein.

3.1.5.1. IS-ACHIEVED-BY **narrative summaries**
Convert argmaps into easy-to-follow narrative summaries that make it easier for evaluators to see the key points they need to.

Perhaps we can do so taking advantage of rhetorical structure theory?

3.1.6. HAS-PART **complete argumentation**
i.e. the evaluation includes a comprehensive overview of the arguments for and against each proposed solution idea

3.1.6.1. IS-VIOLATED-BY **missing arguments**
An idea is missing some important arguments for or against it.

3.1.6.1.1. IS-CAUSED-BY **self-focused**
The participants generate arguments that refer to criteria that concern them personally, but not those that impact welfare for other groups.

3.1.6.1.2. IS-HANDLED-BY **neglected criteria**
Few/no arguments have been created that evaluate a given idea with respect to one of the solution criteria for that problem.

3.1.6.1.3. IS-HANDLED-BY **few/no arguments**
i.e. there are (important) ideas that have few or no arguments attached to them

3.1.6.1.4. IS-HANDLED-BY **unbalanced arguments**
There is a large imbalance in the number of pros and cons in the debate over an idea.

3.1.6.1.5. IS-HANDLED-BY **few people contributed arguments**

3.1.6.1.6. IS-HANDLED-BY **idea/argument rating disconnect**
We can use such techniques as Bayesian inference (Bolstad, 2010) to propagate a user's ratings for arguments up the argument map to predict how the user should have rated the ideas these arguments address. If there is a large divergence between a user's predicted and actual ratings for an idea, that suggests that the user has not yet entered arguments that are compelling to him or her. A suggestion is to show users the possible misalignment between popularity ratings assigned freely by users and computed scores based on the structure of arguments and how much underlying pros and cons are supported. The gap can be used also as an incentive to users: a popular idea with poor computed score should invite its supporters to provide new arguments or improve the existing ones.

3.1.7. HAS-PART **high-quality argumentation**
The arguments entered in the summary map are well-founded.

3.1.7.1. IS-VIOLATED-BY **false premises**
the arguments made are based on false premises

3.1.7.1.1. IS-CAUSED-BY **argument sabotage**
Someone who disagrees with an argument edits it to sabotage it, rather than (properly) simply downrating or arguing against it.

3.1.7.1.1.1. IS-HANDLED-BY **argument edit wars**
assess the prevalence of edit wars (rapid alternating rollbacks) in the post edit histories – esp. by people who take differing positions on the issues they are editing arguments for. see ZIF workshop paper by János Kertész, Budapest: Edit wars on the Wikipedia:. an interesting measure for wikipedia article controversiality = # mutual reverts by more established user accounts (discounting young vandals)

3.1.7.1.2. IS-HANDLED-BY **low argument ratings**
An argument got a low average rating from the community.

3.1.7.1.3. IS-HANDLED-BY **expert evaluation**
Expert(s) judged that the argument is ill-founded.

3.1.7.2. IS-VIOLATED-BY **incorrect inference**
the arguments made are based on logical fallacies

3.1.7.2.1. IS-HANDLED-BY **automated feedback**
The use of artificial intelligence techniques holds promise to increase the effectiveness of argumentation systems by automatically analyzing user actions and providing supportive feedback. see: http//www.ascilite.org.au/ajet/ajet25/butchart.pdf Mclaren, B. M., Scheuer, O., & Mikšátko, J. (2010). Supporting collaborative learning and e-Discussions using artificial intelligence techniques. International Journal of Artificial Intelligence in Education, 20(1)(1), 1-46. Scheuer, O., McLaren, B. M., Loll, F., & Pinkwart, N. (2012). Automated Analysis and Feedback Techniques to Support Argumentation: A Survey. In: N. Pinkwart, & B. M. McLaren (Eds.), Educational Technologies for Teaching Argumentation Skills (pp. 71–124). Bentham Science Publishers.

3.1.7.2.2. IS-HANDLED-BY **expert judgment**
Expert(s) judge that the inference is faulty.

3.1.8. HAS-PART **high-quality ratings**
The community ratings for the argument map contents (i.e. whether issues are relevant, criteria are important, ideas are promising, arguments are compelling) are accurate.

3.1.8.1. IS-VIOLATED-BY **too few ratings**
There are too few ratings for a post to draw reliable conclusions about how the community judges it's value.

3.1.8.2. IS-VIOLATED-BY **dishonest ratings**
Ratings are dishonest.

3.1.8.2.1. IS-HANDLED-BY **rating inconsistency**
A participant gives inconsistent ratings e.g. they rate the arguments supporting an idea highly, but give the idea itself a poor rating.

3.1.8.3. IS-VIOLATED-BY **incorrect ratings**
the user's ratings for the posts are incorrect

3.1.8.3.1. IS-HANDLED-BY **missing suppoprt**
someone gave a strong + or - rating without a backup argument that they authored or highly rated. This can take several levels, according to the Discourse Quality Index (Jurgen Steiner): (1) The speaker does not present any arguments (asks, for

example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

3.1.8.3.2. IS-HANDLED-BY **ignored arguments**
User did not attend to (read, rate) arguments when rating posts impacted by these arguments

3.1.8.3.3. IS-HANDLED-BY **irrational ratings**
see how well a model of rational rating predicts user's ratings. We can use knowledge about map structure to infer what a user's rating for a post "should" be given their ratings for posts below it in tree. For example, with mutex ideas, a rate for X is probably a rate against alternative to X.

3.2. REQUIRES **complete evaluation**
All the (promising) ideas are evaluated.

3.2.1. IS-VIOLATED-BY **evaluation groupthink**
everybody quickly converges to evaluating a very small set of ideas for an issue, ignoring the rest

3.2.1.1. IS-HANDLED-BY **attention narrowing**
we can measure when one idea under an issue is receiving the bulk of the community's argumentation and rating while competing ideas are neglected

4. HAS-PART **P: select best solution**

4.1. REQUIRES **outcome is broadly accepted**

4.1.1. IS-VIOLATED-BY **divisive issues**

4.1.1.1. IS-HANDLED-BY **highlight cross-cutting arguments**
Highlight existing arguments that appeal across balkanized groups in order to help develop increased consensus.

4.1.1.2. IS-HANDLED-BY **high pro/con activity**
do a histogram of activity level for different post types and see if pros and cons are unusually frequent

4.1.1.3. IS-HANDLED-BY **balkanizing issues**
we can use latent semantic indexing or principal components analysis or the like to find which are the issue sets that most divide people into clusters. A principal components analysis could help find, in effect, the fault lines in a debate, the sets of issues that most tend to divide people. We could find, for example, that abortion and gun control and school vouchers are highly divisive issues but if people agree on one of those issues they tend to agree on all the others well. We can then ask: what do these issues have in common? What underlying motivation or belief do they reflect? How can we attempt to reduce polarization along this dimension?

4.1.1.4. IS-HANDLED-BY **issues with high/growing rating variance**
where there are many arguments and contributors but no clear preponderance of highly-rated pros or cons

4.1.2. IS-VIOLATED-BY **many disaffected participants**
There are many deliberation participants who feel the selected outcome is unacceptable.

4.1.2.1. IS-HANDLED-BY **P: resolve: identify commonalities among participants**
cf Terry Steichen's work (the TopicCentral system) on finding commonalities in different people's favored portions of the deliberation map.

4.1.2.2. IS-HANDLED-BY **P: resolve: engage conflict resolution experts**
identified perhaps using analytics applied to deliberation summary?

4.2. REQUIRES **solution map is mature**
i.e. there is sufficient coverage of the issues, ideas, and arguments to make a decision

4.2.1. IS-ACHIEVED-BY **author/moderator activity dropoff**
If the fraction of author vs moderator contributions to a discussion drops, this suggests that the discussion is losing steam - it is only kept active by the effort of the moderators.

4.2.2. IS-ACHIEVED-BY **six hats**
Assess whether the problem solving session has progressed through a complete \"six hats\" program: Blue, White, Green, Red, Yellow, Black, which can be mapped to an argument map setting as follows:

- Thinking (Blue) - thinking about thinking, process issues
- Information: (White) - considering what information is available, what are the facts? provide details on the issues (= problems to be solved)
- Creativity (Green) - statements of provocation and investigation, seeing where a thought goes (= propose ideas)
- Good points judgment (Yellow) - logic applied to identifying benefits, seeking harmony (= identifying pros)
- Bad points judgment (Black) - logic applied to identifying flaws or barriers, seeking mismatch (= identifying cons)
- Emotions (Red) - instinctive gut reaction or statements of emotional feeling (but not any justification) (=> ratings? comments?) indicative of final stage - detected using prevalence of emotive words?

wikipedia article

4.2.3. IS-ACHIEVED-BY **"full" map topology**
map is both sufficiently bushy and deep.

4.2.4. IS-ACHIEVED-BY **completes narrative template**
The customer specifies the kind of narrative they want i.e. the main questions, the depth of argumentation, the breadth of options etc The system evaluates how far the argument map has gone to enabling that narrative, and asks the crowd to focus on the areas that yet need to be filled in. See work on rhetorical structures e.g.

http//www.cs.columbia.edu/~kathy/NLP/ClassSlides/Slides09/Class20-Discourse/my-discourse.pdf

4.2.5. IS-ACHIEVED-BY **lifecycle stages**
There are many different possible models of the life stages a deliberation goes through as it matures. These include: • evolve from defining issues to proposing ideas to identifying increasingly broad and deep trees of pro and con arguments • evolve from creating new posts, to refining them, followed eventually by relative quiescence • opinion churn (i.e. whether the highest-rated ideas for individuals, as well as the community as a whole, are still changing rapidly or not) moderates as we reach the end of the lifecycle. • community support (as assessed by idea and arg ratings) concentrates on a few strongly supported ideas (lots of high ratings) • deliberation goes through the stages of preach to crowd, angry debunkers, filling in implicit support with reasoned data-based responses, irrelevant bored commentaryF • map growth tends to follow an S-shaped curve: map may be reaching maturity when slope decreases. http//crowdresearch.org/blog/?p=4602&utm_source=feedburner&utm_medium=email&utm_campaign=Feed%3A+FollowTheCrowd+%28Follow+the+Crowd%29

4.3. REQUIRES **high-quality votes**
i.e. they reflect the user's best judgment about which selection to make

4.3.1. HAS-PART **votes are truthful**

4.3.2. HAS-PART **votes are rational**
i.e. they represent a logically consistent response to all the relevant ideas and arguments (as opposed to some kind of bias)

4.3.2.1. IS-VIOLATED-BY **ignore higher-level context**
Users make decisions for low-level issues with taking into account higher-level issues and decisions that should have a major impact.

4.3.2.1.1. IS-HANDLED-BY **avoid: encourage hierarchical rating**
vote/rate on abstraction before voting on details

4.3.2.2. IS-VIOLATED-BY **hedgehog voter**
The voter has ignored ideas or arguments that should be relevant to the issue they voted on. We can call this the "hedgehog" exception, after Philip Tetlock, who pointed out that some people ("hedgehogs") only pay attention to a subset of information that is close to their original point of view (see Zaller), while other roam the information space more broadly ("foxes").

4.3.2.2.1. IS-HANDLED-BY **opinion shift**
If a user has substantially changed his ratings about ideas and arguments, this suggests he/she is open to being influenced by new information and perspectives.

4.3.2.2.2. IS-HANDLED-BY **user saw relevant posts**
Check if user has not read or rated relevant ideas and arguments under the issue he/she is voting on.

4.3.2.3. IS-VIOLATED-BY **voting cascades**
It has been shown that when people are asked to rate competing ideas, if they can see the ratings made to date (e.g. they see the ideas in popularity-sorted order), then the first ideas that happen to get a rating advantage tend to become the eventual winners—they “lock in” to the winning position—even if they are worse than ideas that appeared later or started with lower ratings (Salganik et al., 2006). its is therefore a problem if people vote for ideas based on their popularity (i.e. based on how many other people have voted for them) rather than their inherent merits.

4.3.2.3.1. IS-HANDLED-BY **ratings lock**
check whether the popularity order for a set of competing ideas remains relatively unchanged as the deliberation progresses

4.3.2.4. IS-VIOLATED-BY **bias**
participants is biased towards a given decision irregardless of arguments and other alternatives

4.3.2.4.1. IS-HANDLED-BY **motivated position change**
We can assess degree of bas by measuring whether users change their position through the course of the deliberation, or not. This can have several levels, according to Jurgen Steiner's Discourse Quality Index: (1) The speaker indicates a change of position. Gives as reason for change arguments heard during the experiment. (2) The speaker indicates a change of position. Does not refer to arguments heard during the experiment. (3) The speaker does not indicate a change in position. But does acknowledge the value of other positions heard during the experiment. (4) The speaker does not indicate a change of position. And does not acknowledge the value of other positions heard during the experiment.

4.3.2.4.2. IS-HANDLED-BY **coherence theory**
Can we use coherence theory, applied to the posts that the user ranked highly and thus presumably used in their decision, to assess the logical coherence of their votes? See http://www.iiia.csic.es/~joseph/index2.html

4.3.2.4.3. IS-HANDLED-BY **rating disconnect**
Assuming all the key arguments have be entered i.e. the map is mature: the user selects an idea that doesn't make sense given the ratings he/she gave to their underlying arguments and the competing ideas.

4.3.3. HAS-PART **votes are truthful**

4.3.4. HAS-PART **votes are rational**
i.e. they represent a logically consistent response to all the relevant ideas and arguments (as opposed to some kind of bias)

4.3.4.1. IS-VIOLATED-BY **ignore higher-level context**
Users make decisions for low-level issues with taking into account higher-level issues and decisions that should have a major impact.

4.3.4.1.1. IS-HANDLED-BY **avoid: encourage hierarchical rating**
vote/rate on abstraction before voting on details

4.3.4.2. IS-VIOLATED-BY **hedgehog voter**
The voter has ignored ideas or arguments that should be relevant to the issue they voted on. We can call this the "hedgehog" exception, after Philip Tetlock, who pointed out that some people ("hedgehogs") only pay attention to a subset of information that is close to their original point of view (see Zaller), while other roam the information space more broadly ("foxes").

4.3.4.2.1. IS-HANDLED-BY **opinion shift**
If a user has substantially changed his ratings about ideas and arguments, this suggests he/she is open to being influenced by

new information and perspectives.

4.3.4.2.2. IS-HANDLED-BY **user saw relevant posts**
Check if user has not read or rated relevant ideas and arguments under the issue he/she is voting on.

4.3.4.3. IS-VIOLATED-BY **voting cascades**
It has been shown that when people are asked to rate competing ideas, if they can see the ratings made to date (e.g. they see the ideas in popularity-sorted order), then the first ideas that happen to get a rating advantage tend to become the eventual winners—they "lock in" to the winning position—even if they are worse than ideas that appeared later or started with lower ratings (Salganik et al., 2006). its is therefore a problem if people vote for ideas based on their popularity (i.e. based on how many other people have voted for them) rather than their inherent merits.

4.3.4.3.1. IS-HANDLED-BY **ratings lock**
check whether the popularity order for a set of competing ideas remains relatively unchanged as the deliberation progresses

4.3.4.4. IS-VIOLATED-BY **bias**
participants is biased towards a given decision irregardless of arguments and other alternatives

4.3.4.4.1. IS-HANDLED-BY **motivated position change**
We can assess degree of bas by measuring whether users change their position through the course of the deliberation, or not. This can have several levels, according to Jurgen Steiner's Discourse Quality Index: (1) The speaker indicates a change of position. Gives as reason for change arguments heard during the experiment. (2) The speaker indicates a change of position. Does not refer to arguments heard during the experiment. (3) The speaker does not indicate a change in position. But does acknowledge the value of other positions heard during the experiment. (4) The speaker does not indicate a change of position. And does not acknowledge the value of other positions heard during the experiment.

4.3.4.4.2. IS-HANDLED-BY **coherence theory**
Can we use coherence theory, applied to the posts that the user ranked highly and thus presumably used in their decision, to assess the logical coherence of their votes? See http://www.iiia.csic.es/~joseph/index2.html

4.3.4.4.3. IS-HANDLED-BY **rating disconnect**
Assuming all the key arguments have be entered i.e. the map is mature: the user selects an idea that doesn't make sense given the ratings he/she gave to their underlying arguments and the competing ideas.

4.4. REQUIRES **sufficient votes**
sufficient votes are available to fully, and fairly, capture the wisdom and preferences of the voters.

4.4.1. IS-VIOLATED-BY **insufficient votes**
There is insufficient preference information (in terms of votes or ratings) to pick a clear winner among the solution ideas.

4.4.1.1. IS-HANDLED-BY **confidence analysis**
Can we do an analysis to determine which ideas need to be assessed more completely in order to allow high-confidence selections of top-level solution ideas? Would we need some kind of confidence interval analysis? Could this be calculated based just on simple ratings, or would people need to express confidence scores for their ratings (e.g. very sure, not very sure).It would make sense to take into account the controversiality of the ideas e.g. if an idea is controversial, we would probably want to get more ratings for it to be more sure that people really prefer it (or not).

4.5. REQUIRES **votes aggregated properly**

4.5.1. HAS-PART **representative**
i.e. the decision created by aggregating the group's vote should represent what most individuals wanted

4.5.2. HAS-PART **fair**

4.5.3. HAS-PART **representative**
i.e. the decision created by aggregating the group's vote should represent what most individuals wanted

4.5.4. HAS-PART **fair**

5. HAS-PART **Enact solution**

5.1. REQUIRES **feasible solution**
The solution is a feasible one (i.e. can be implemented).

5.2. REQUIRES **commitment to action**
In social media, participation is high but incitement to action is historically low. Online debate and deliberation tools are populate by enthusiasts who have interest in the subject, spend time and efforts into debating it, but have not yet committed into taking action. How do we engage enthusiast/motivated audiences to translate the emerging trends and patterns into concrete actions to lead to further change?

6. REQUIRES **good process**

6.1. HAS-PART **TBD**
This is where we can attach new metrics that we haven't placed in the model yet.

6.2. HAS-PART **the right participants are involved**
i.e. people with the necessary depth and diversity of perspectives and skills

6.2.1. IS-VIOLATED-BY **too few authors**
The ideas for an issue come from an especially small number of contributors

6.2.1.1. IS-CAUSED-BY **newbie attrition**
newbies are discouraged by early edits being reverted/uncertified see: http://crowdresearch.org/blog/?p=1907

6.2.1.1.1. IS-HANDLED-BY **short-lived activity**
i.e. a user participates actively for a short while after joining, then stops for a prolonged period.

6.2.1.1.2. IS-HANDLED-BY **encourage newbies**
encourage and retain new users ee: http://crowdresearch.org/blog/?p=1907

6.2.1.2. IS-HANDLED-BY **gini coefficient**

Gini coefficients range between 0 and 1: 0 → perfect equality (all participants contributing the same number of posts) 1 → perfect inequality (one participant contributing all posts and everyone else contributing none).

6.2.1.3. IS-HANDLED-BY **narrow contributions histogram**
If we plot the activity of each user as a bar plot, sorted left to right by activity, we can assess what proportion of the users are active or not. A narrow peak of high activity implies few people are active.

6.2.2. IS-VIOLATED-BY **inadequate author diversity**
ideas come from "non-disjoint" author sets i.e. where all the authors tend to agree about the pros and cons for different ideas and therefore probably share an intellectual frame

6.2.2.1. IS-HANDLED-BY **ask participants to suggest new members**
Ask contributors to suggest people with alternative views: "fresh blood" for the deliberation.

6.2.2.2. IS-HANDLED-BY **authors have similar rating vectors**
We can perform vector orthogonalization (Householder, 1958) on authors' rating vectors, followed by a simple vector distance calculation, to assess how much the opinions for different authors diverge.

6.2.2.3. IS-HANDLED-BY **show contributor demographics**
If demographics are available, we can check for diversity, or lack of it, in the participant population.

6.2.2.4. IS-HANDLED-BY **encourage participation from underrepresented demographics**

6.3. HAS-PART **participants have good inputs**
i.e. they are exposed to a diverse range of materials to inform their ideation and decision making

6.3.1. IS-VIOLATED-BY **myopic authoring**
Authors devote themselves to building upon their own contributions without also refining/critiquing content contributed by others.

6.3.1.1. IS-HANDLED-BY **self/other ratio**
measure (N°rating + N° Pro/Con as answer to other's posts)/ N° of Own Post

6.3.2. IS-VIOLATED-BY **miss relevant content**
A user misses content that would elicit more contributions from them, if they had seen it.

6.3.2.1. IS-HANDLED-BY **renewed interest**
Telligent folks (Marc Smith et al) showed that some lurkers (10-20%) would contribute further if they knew that posts that interested them (they had viewed edited rated commented on them) had become "hot" and therefore worth spending time on, so it's good if a system can notify people when that happens

6.3.2.2. IS-HANDLED-BY **viewed by people w/similar interests**
Use some kind of clustering (e.g. based on vector de-orthogonalization) to find people who have interests like me, and notify me about activity (views, edits, comments, rates, hots) that was interesting to them. do an eigenvector analysis on rating vectors (perhaps we can augment this by taking advantage of tree structure of map?) and look for what was viewed by people located near to you in that eigenspace Maybe one way to look at the problem is that we want to define, for each person, an estimate of the likelihood that they like/dislike each idea. This suggests two thoughts: () people may add both pros and cons for an idea: how do we combine these? Perhaps we need *two* scores for each idea: one p(like), the other p(dislike)? Or we simply say p(like) = 0.5? () if we assign each person a vector that gives the probability that they like each idea, then the default value can be 0.5 (equal chance of liking it or not), so there are no missing values, simplifying the user similarity calculation. We can start with very simple rules for setting the p(like) values, and refine as needed. () if we propagate p(like/dislike) values up the deliberation map, which does seem potentially useful, we may need to use an evidence accumulation math e.g. Bayesian. For example, if I have two separate lines of evidence for believing conclusion1, the p(conclusion1) is given by the min or max (and not the sum) of P(evidence1) and P(evidence2), depending on whether the evidence pieces have an AND or OR relationship. Currently, the default semantics for all arguments is OR, though I've considered adding AND nodes to allow correct propagation of belief values up the argument map. () how do we propagate p(like) values up past issues? In particular, should we take the issue rating (which is intended to capture the user's estimate of the importance of the issue) into account? I may really like an idea for an unimportant issue, for example: how much impact should that have up the tree?

6.3.2.3. IS-HANDLED-BY **favored post with low/declining ratings**
Lets author know if he/she should try to add arguments that move the community support in the direction I want

6.3.2.4. IS-HANDLED-BY **busy topics**
calculate an activity score, point people towards busy stuff to be relative to other branches, or number of users, or …? (to avoid penning out scores) to be aggregated up tree? Simply keep track of current score plus time since last update – update value whenever a new event occurs, or when read – events propagate up but reduce in strength as they go - make decay rate faster, so gives a current picture of activity

6.3.2.5. IS-HANDLED-BY **near areas that interested me**
Notify user of activity on posts that are "nearby" to posts they have been interested in (viewed, edited, created, rated, hotted) in the past.

6.3.2.6. IS-HANDLED-BY **interesting to people in my social network**
we can notify users of when there is post activity (views, edits, comments, rates, hots) either initiated by, or considered to be interesting to, people in the user's social network

6.3.3. IS-VIOLATED-BY **static subgroups**
(small and) relatively static sets of people work on each part of the deliberation, so there is little "fresh blood", new ideas, new perspectives

6.3.3.1. IS-HANDLED-BY **encourage people to shift topics**
... to break up static groups

6.3.3.2. IS-HANDLED-BY **track subgroups**
... to see whether a small static group has consistently been responsible for all the content in a given topic area in the argument map.

6.3.4. IS-VIOLATED-BY **platform islands**

Community participants use different tools to support online debate and conversations then remain locked within tools. This implies that topics, ideas and outcomes of online conversations remain constrained to specific communities and fail to cross-federate debate across platforms.

6.3.4.1. IS-HANDLED-BY **contributions clustered by platform**
We can compare the argument maps for different platforms to see whether or not contributions are clustered by platform, or not, i.e. whether key content appears in just one or a few of the platforms.

6.3.4.2. IS-HANDLED-BY **"seeders" propagate ideas cross-platform**
"seeders" transfer key ideas from the argument map summary to social media platforms where they did not previously appear.

6.3.5. IS-VIOLATED-BY **balkanization**
“balkanization” means: the community self-organizes into cliques that agree within themselves but disagree with each other. It occurs when a community divides itself into partisan sub-groups where members of each group agree with one other but actively fight against groups with competing ideas. This can be a problem if it means that the sub-groups do not build upon potentially valuable ideas from other groups because of in-group/out-group social dynamics. Cliques form wherein each clique is devoted to a particular class of solutions and either ignores or actively argues against all other ideas, rather than seeing whether new ideas can be created that combine the best features of both.

6.3.5.1. IS-HANDLED-BY **insularity**
There are multiple subgroups (defined by social network analysis) which discuss related topics but do not talk to each other.

6.3.5.2. IS-HANDLED-BY **odd couples**
Bring an author's attention to ideas that come from an author whose interest/rating vector is very *different*, thus fighting balkanization.

6.3.5.3. IS-HANDLED-BY **bias detectors**
Tools are now emerging to detect whether people are one-sided in their news reading (e.g. see http://crowdresearch.org/blog/?p=8244). Perhaps these can be adapted to detect when innovation contributors are one-sided in the inputs they are taking in.

6.3.5.4. IS-HANDLED-BY **attitude space clusters**
for each user, calculate their attitudes towards an idea (i.e. refining idea or uprating or adding pro arguments is positive, downrating or adding con arguments is negative) and (perhaps after singular vector decomposition) look for clusters of distinct groups that are similar within but very different from each other).

6.3.6. IS-ACHIEVED-BY **harvest new content**
This is the process wherein harvesters scan through social media in order to find issues, ideas, arguments that can contribute to the social innovation engagement.

6.3.6.1. HAS-PART **summarize as map**
Organize content the authors found into a summary argument map.

6.3.6.1.1. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.3.6.1.1.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.3.6.1.1.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.3.6.1.1.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.3.6.1.2. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.3.6.1.2.1. REQUIRES **find correct tag**

6.3.6.1.2.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.3.6.1.2.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.3.6.1.3. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.3.6.1.3.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.3.6.1.4. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.3.6.1.4.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.3.6.1.4.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.3.6.1.4.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.3.6.1.5. HAS-PART **tag atoms**

Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.3.6.1.5.1. REQUIRES **find correct tag**

6.3.6.1.5.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.3.6.1.5.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.3.6.1.6. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.3.6.1.6.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.3.6.2. REQUIRES **find useful content**
Authors search (e.g. social media) to find (all) content relevant to the social innovation engagement.

6.3.6.2.1. IS-ACHIEVED-BY **sentiment analysis**
use sentiment analysis to troll web to help harvesters find controversy - e.g. negative sentiment probably means a con.

6.3.6.2.2. IS-ACHIEVED-BY **weblinks from fertile sources**
We can mine web link structure to suggest sources e.g. an article section that proposed a pro argument may have links to other pages which are probably cited to support the pro argument

6.3.6.2.3. IS-ACHIEVED-BY **tagged fertile sources**
harvesters can tag social media site pages as fertile or not, so others can also benefit from that resource.

6.3.6.2.4. IS-ACHIEVED-BY **used in argmap**
use argmap post backlinks to suggest fertile social media sources. If a site had good info, go back to see if more is available

6.3.6.2.5. IS-ACHIEVED-BY **find important authors via SNA**
social network analytics (centrality measures and community detection algorithms) can be used to identify highly influential individuals and groups whose outputs may be particularly worthy of harvesting.

6.3.6.2.6. IS-ACHIEVED-BY **topic trends**
Use a trend detection tool such as that provided by google to find hot topics that may be ripe for harvesting.

6.3.6.2.7. IS-ACHIEVED-BY **forwarding statistics**
forwarding relationships in email and microblogs such as twitter can be used to detect possible atoms of interest. Text that is frequently re-forwarded/quoted, for example, might be especially worthy of a harvester's attention

6.3.6.3. HAS-PART **summarize as map**
Organize content the authors found into a summary argument map.

6.3.6.3.1. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.3.6.3.1.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.3.6.3.1.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.3.6.3.1.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.3.6.3.2. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.3.6.3.2.1. REQUIRES **find correct tag**

6.3.6.3.2.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.3.6.3.2.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.3.6.3.3. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.3.6.3.3.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.3.6.3.4. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.3.6.3.4.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.3.6.3.4.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of

interest, so they can be merged.

6.3.6.3.4.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.3.6.3.5. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.3.6.3.5.1. REQUIRES **find correct tag**

6.3.6.3.5.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.3.6.3.5.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.3.6.3.6. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.3.6.3.6.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.4. HAS-PART **participants contribute fully**
participants contribute fully in terms of their time and skills

6.4.1. HAS-PART **critical mass**
There are enough active contributors to create self-sustaining deliberation.

6.4.1.1. IS-ACHIEVED-BY **author/moderator effort ratio**
If moderators account for a disproportionately high fraction of the overall interactions, that suggests that there are not enough participants to achieve a self-sustaining deliberation.

6.4.2. HAS-PART **retain productive contirbutors**
Keep productive contributors involved.

6.4.3. HAS-PART **strong incentives for participation**

6.4.3.1. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.4.3.1.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.4.3.2. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.4.3.2.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.4.3.3. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.4.3.3.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.4.3.4. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.4.3.4.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.4.3.5. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.4.3.5.1. IS-ACHIEVED-BY **contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.4.3.6. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.4.3.6.1. IS-ACHIEVED-BY **use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.4.3.7. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.4.3.7.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.4.3.8. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.4.3.8.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.4.3.9. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.4.3.9.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.4.3.10. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.4.3.10.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.4.3.11. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.4.3.11.1. IS-ACHIEVED-BY **contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.4.3.12. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.4.3.12.1. IS-ACHIEVED-BY **use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.4.4. HAS-PART **few disincentives for participation**

6.4.4.1. IS-VIOLATED-BY **abusive behavior**
Abusive behavior by participants can be a disincentive for participation by other participants.

6.4.4.1.1. IS-HANDLED-BY **foul/respectful language**
We can search for language that reveals respectfulness of community interactions e.g. The speaker uses respectful language towards other participants and/or their arguments. Include also moderately respectful language, not only statements such as “your argument is truly brilliant” but also statements such as “your argument is not bad.” The speaker uses foul language to attack other participants on a personal level. Include also mild foul language, not only statements such as “you are a liar” but also statements such as “you seem a little confused.” Code the names of the participants attacked in this way and give the exact quote of the foul language. The speaker uses foul language to attack the arguments of other participants but abstains from personal attacks. Here again include also mild foul language, not only statements such as “this argument is stupid” but also statements such as “this argument is a little weak.” Code the names of the participants whose arguments are attacked in this way and give the exact quote of the foul language.

6.4.5. IS-VIOLATED-BY **one-sided contribution**
participants focus on just a single style of contribution (e.g. adding ideas, or adding arguments) and thus potentially are not contribute some of their skills to the deliberation

6.4.5.1. IS-CAUSED-BY **idea nay-sayer**
User has only critiques, no positive suggestions.

6.4.5.1.1. IS-HANDLED-BY **only cons for issue ideas**
user doesn’t like any of existing options for an issue

6.4.5.1.2. IS-HANDLED-BY **suggest alternative**
ask user to suggest a new alternative for an issue

6.4.5.2. IS-HANDLED-BY **six hats stats**
Assess whether the user has used all "six hats":

- Thinking (Blue) - thinking about thinking, process issues
- Information: (White) - considering what information is available, what are the facts? provide details on the issues (= problems to be solved)
- Creativity (Green) - statements of provocation and investigation, seeing where a thought goes (= propose ideas)
- Good points judgment (Yellow) - logic applied to identifying benefits, seeking harmony (= identifying pros)
- Bad points judgment (Black) - logic applied to identifying flaws or barriers, seeking mismatch (= identifying cons)
- Emotions (Red) - instinctive gut reaction or statements of emotional feeling (but not any justification) (=> ratings? comments?) indicative of final stage - detected using prevalence of emotive words?

wikipedia article

6.4.5.3. IS-HANDLED-BY **unjustified contributions**
The participant provides ideas and arguments without providing arguments backing them up. This can have several levels (according to Jurg Steiner's Discourse Quality Index): (1) The speaker does not present any arguments (asks, for example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

6.4.5.4. IS-HANDLED-BY **contributions histogram**

Create histogram assessing user contributions for post types (ideas, issues, arguments) and topics (branches of the map) to detect if they are specializing in a narrow scope.

6.4.6. IS-VIOLATED-BY **non-participation**
Community members are not participating in the deliberation.

6.4.6.1. IS-HANDLED-BY **lapsed contributors**
Detect people who started but stopped participating.

6.4.6.2. IS-HANDLED-BY **user activity stats**

6.4.7. IS-ACHIEVED-BY **calculate contribution scores**
you can add points to capacity building exercises, like if you read x or watch the video or do the training well you get points they assign points to you if you operate in the directions that the organizers wants, so for example when you do a post right you get a point. when you do a mistake you lose one for creation and evaluation and moderation and prediction accuracy You can be offered personalized point-gaining opportunities, identified by the tutoring and attention mediation heuristics i.e. metrics that identify how much a user has contributed. Some possible sources of ideas for how to calculate these scores include: • slashdot (karma points) • digg • yourview.org.au (credibility score) http://www.publicdeliberation.net/jpd/vol8/iss1/art12/ be a thoughtful and constructive participant - particularly in the eyes of other high-credibility participants, and particularly in relation to those who disagree with you. Participate on a wide variety of issues; Participate fully, i.e. by rating, commenting and voting Give due consideration to the relevant arguments on both sides Earn the respect of others, particularly those who already have high credibility, and those who disagree with you. Don't be abusive, rude, obscene, arrogant, or obnoxious.

6.4.7.1. REQUIRES **accurate scores**
The contribution scores should be accurate.

6.4.7.1.1. IS-VIOLATED-BY **reputation gaming**
Users can try to game the system to get high reputation scores without actually contributing much to the social innovation process.

6.4.8. HAS-PART **critical mass**
There are enough active contributors to create self-sustaining deliberation.

6.4.8.1. IS-ACHIEVED-BY **author/moderator effort ratio**
If moderators account for a disproportionately high fraction of the overall interactions, that suggests that there are not enough participants to achieve a self-sustaining deliberation.

6.4.9. HAS-PART **retain productive contirbutors**
Keep productive contributors involved.

6.4.10. HAS-PART **strong incentives for participation**

6.4.10.1. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.4.10.1.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.4.10.2. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.4.10.2.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.4.10.3. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.4.10.3.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.4.10.4. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.4.10.4.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.4.10.5. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.4.10.5.1. IS-ACHIEVED-BY **contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.4.10.6. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.4.10.6.1. IS-ACHIEVED-BY **use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.4.10.7. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.4.10.7.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.4.10.8. HAS-PART **reputation**

High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.4.10.8.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.4.10.9. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.4.10.9.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.4.10.10. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.4.10.10.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.4.10.11. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.4.10.11.1. IS-ACHIEVED-BY **:contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.4.10.12. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.4.10.12.1. IS-ACHIEVED-BY **:use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.4.11. HAS-PART **few disincentives for participation**

6.4.11.1. IS-VIOLATED-BY **abusive behavior**
Abusive behavior by participants can be a disincentive for participation by other participants.

6.4.11.1.1. IS-HANDLED-BY **foul/respectful language**
We can search for language that reveals respectfulness of community interactions e.g. The speaker uses respectful language towards other participants and/or their arguments. Include also moderately respectful language, not only statements such as “your argument is truly brilliant” but also statements such as “your argument is not bad.” The speaker uses foul language to attack other participants on a personal level. Include also mild foul language, not only statements such as “you are a liar” but also statements such as “you seem a little confused.” Code the names of the participants attacked in this way and give the exact quote of the foul language. The speaker uses foul language to attack the arguments of other participants but abstains from personal attacks. Here again include also mild foul language, not only statements such as “this argument is stupid” but also statements such as “this argument is a little weak.” Code the names of the participants whose arguments are attacked in this way and give the exact quote of the foul language.

6.5. HAS-PART **participants contribute effectively**

6.5.1. IS-VIOLATED-BY **troublemakers**
Troublemaker users are reducing the effectiveness of the social innovation system.

6.5.1.1. IS-HANDLED-BY **count trouble tags**
Count trouble tags created in response to actions by that user.

6.5.2. IS-VIOLATED-BY **participants don't know where to contribute**
Participants don't know what parts of the social innovation engagement they can best contribute to.

6.5.2.1. IS-HANDLED-BY **:attention mediation**
The system notifies participants about tasks that need attention and that they are suited to perform.

6.5.2.1.1. HAS-PART **:gather deliberation data**
on both the users and the content they generate.

6.5.2.1.1.1. REQUIRES **sufficient data is available**

6.5.2.1.1.1.1. IS-VIOLATED-BY **insufficient deliberation info**

6.5.2.1.1.1.1.1. IS-HANDLED-BY **:active learning**
Use information theory to determine which actions will have the greater positive impact in terms of identifying the most important exceptions. Thanks to Avi for the idea.

6.5.2.1.2. HAS-PART **:run metrics**
run metrics processes to detect symptoms of exceptions

6.5.2.1.2.1. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.5.2.1.2.2. HAS-PART **:mathematical analysis**

6.5.2.1.2.2.1. HAS-PART **:belief propagation**

6.5.2.1.2.2.2. HAS-PART **:network analysis**

6.5.2.1.2.2.3. HAS-PART **:eigenvector analysis**

6.5.2.1.2.2.4. HAS-PART **:belief propagation**

6.5.2.1.2.2.5. HAS-PART **:network analysis**

6.5.2.1.2.2.6. HAS-PART **:eigenvector analysis**

6.5.2.1.2.3. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.5.2.1.2.4. HAS-PART **:mathematical analysis**

6.5.2.1.2.4.1. HAS-PART **:belief propagation**

6.5.2.1.2.4.2. HAS-PART **:network analysis**

6.5.2.1.2.4.3. HAS-PART **:eigenvector analysis**

6.5.2.1.2.4.4. HAS-PART **:belief propagation**

6.5.2.1.2.4.5. HAS-PART **:network analysis**

6.5.2.1.2.4.6. HAS-PART **:eigenvector analysis**

6.5.2.1.3. HAS-PART **:diagnose exceptions**
Determine which exceptions are taking place, given the current metrics values.

6.5.2.1.4. HAS-PART **:prioritize exceptions**
Since exceptions are identified wrt an "ideal" deliberation process, there are likely to be many possible ones in play , more perhaps than can be handled at one time. For this reason, we need a way to prioritize which exceptions receive more attention.

6.5.2.1.5. HAS-PART **:select handlers**
Select which handler, if there are several available, to enact.

6.5.2.1.5.1. IS-REALIZED-BY **:exception-specific handler**
Pick a handler specific to that exception.

6.5.2.1.5.2. IS-REALIZED-BY **:dynamic incentives**
harvesters can be guided by dynamic pricing whose prizes reflect what customer wants, or what the attention mediation algorithms propose

6.5.2.1.5.3. IS-REALIZED-BY **:notification**
If no more specific handler is available, the default handler is to notify the people who could do something about the exception.

6.5.2.1.6. HAS-PART **:run handlers**
Run the selected handler(s) to address the exception

6.5.2.1.7. HAS-PART **:learn from experience**
i.e. learn how to detect and handle exceptions more effectively over time, based on previous experience.

6.5.2.1.7.1. HAS-PART **:collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.5.2.1.7.2. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.5.2.1.7.3. HAS-PART **:machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.5.2.1.7.4. HAS-PART **:collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.5.2.1.7.5. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.5.2.1.7.6. HAS-PART **:machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.5.2.1.8. HAS-PART **:gather deliberation data**
on both the users and the content they generate.

6.5.2.1.8.1. REQUIRES **sufficient data is available**

6.5.2.1.8.1.1. IS-VIOLATED-BY **insufficient deliberation info**

6.5.2.1.8.1.1.1. IS-HANDLED-BY **:active learning**
Use information theory to determine which actions will have the greater positive impact in terms of identifying the most important exceptions. Thanks to Avi for the idea.

6.5.2.1.9. HAS-PART **:run metrics**
run metrics processes to detect symptoms of exceptions

- 6.5.2.1.9.1. HAS-PART **:triplestore queries**
  i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.
- 6.5.2.1.9.2. HAS-PART **:mathematical analysis**
  - 6.5.2.1.9.2.1. HAS-PART **:belief propagation**
  - 6.5.2.1.9.2.2. HAS-PART **:network analysis**
  - 6.5.2.1.9.2.3. HAS-PART **:eigenvector analysis**
  - 6.5.2.1.9.2.4. HAS-PART **:belief propagation**
  - 6.5.2.1.9.2.5. HAS-PART **:network analysis**
  - 6.5.2.1.9.2.6. HAS-PART **:eigenvector analysis**
- 6.5.2.1.9.3. HAS-PART **:triplestore queries**
  i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.
- 6.5.2.1.9.4. HAS-PART **:mathematical analysis**
  - 6.5.2.1.9.4.1. HAS-PART **:belief propagation**
  - 6.5.2.1.9.4.2. HAS-PART **:network analysis**
  - 6.5.2.1.9.4.3. HAS-PART **:eigenvector analysis**
  - 6.5.2.1.9.4.4. HAS-PART **:belief propagation**
  - 6.5.2.1.9.4.5. HAS-PART **:network analysis**
  - 6.5.2.1.9.4.6. HAS-PART **:eigenvector analysis**

6.5.2.1.10. HAS-PART **:diagnose exceptions**
Determine which exceptions are taking place, given the current metrics values.

6.5.2.1.11. HAS-PART **:prioritize exceptions**
Since exceptions are identified wrt an "ideal" deliberation process, there are likely to be many possible ones in play , more perhaps than can be handled at one time. For this reason, we need a way to prioritize which exceptions receive more attention.

6.5.2.1.12. HAS-PART **:select handlers**
Select which handler, if there are several available, to enact.

- 6.5.2.1.12.1. IS-REALIZED-BY **:exception-specific handler**
  Pick a handler specific to that exception.
- 6.5.2.1.12.2. IS-REALIZED-BY **:dynamic incentives**
  harvesters can be guided by dynamic pricing whose prizes reflect what customer wants, or what the attention mediation algorithms propose
- 6.5.2.1.12.3. IS-REALIZED-BY **:notification**
  If no more specific handler is available, the default handler is to notify the people who could do something about the exception.

6.5.2.1.13. HAS-PART **:run handlers**
Run the selected handler(s) to address the exception

6.5.2.1.14. HAS-PART **:learn from experience**
i.e. learn how to detect and handle exceptions more effectively over time, based on previous experience.

- 6.5.2.1.14.1. HAS-PART **:collect user feedback**
  Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)
- 6.5.2.1.14.2. HAS-PART **human visual pattern recognition**
  We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.
- 6.5.2.1.14.3. HAS-PART **:machine learning**
  Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML
- 6.5.2.1.14.4. HAS-PART **:collect user feedback**
  Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)
- 6.5.2.1.14.5. HAS-PART **human visual pattern recognition**
  We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.
- 6.5.2.1.14.6. HAS-PART **:machine learning**
  Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even

learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.5.3. IS-VIOLATED-BY **users don't understand how to use system**
Users are lacking information or skills needed to use the social innovation system effectively.

6.5.3.1. IS-HANDLED-BY **custom training**
Users can be pointed to help materials that appear relevant to the problems they seem to be having with using the deliberation system correctly. See the reasons that moderators have identified for rejecting posts (e.g. improperly unbundled or located), use this to provide some kind of personalized training for the user - a basis for integrated tutoring system.

6.5.3.2. IS-HANDLED-BY **slow certification**
many moderator iterations are needed for certification of the contributor's posts

6.5.3.3. IS-HANDLED-BY **high action/output ratio**
i.e. users try a lot of edit actions, but produce only a few certified posts as a result

6.6. HAS-PART **TBD**
This is where we can attach new metrics that we haven't placed in the model yet.

6.7. HAS-PART **the right participants are involved**
i.e. people with the necessary depth and diversity of perspectives and skills

6.7.1. IS-VIOLATED-BY **too few authors**
The ideas for an issue come from an especially small number of contributors

6.7.1.1. IS-CAUSED-BY **newbie attrition**
newbies are discouraged by early edits being reverted/uncertified see: http://crowdresearch.org/blog/?p=1907

6.7.1.1.1. IS-HANDLED-BY **short-lived activity**
i.e. a user participates actively for a short while after joining, then stops for a prolonged period.

6.7.1.1.2. IS-HANDLED-BY **encourage newbies**
encourage and retain new users ee: http://crowdresearch.org/blog/?p=1907

6.7.1.2. IS-HANDLED-BY **gini coefficient**
Gini coefficients range between 0 and 1: 0 → perfect equality (all participants contributing the same number of posts) 1 → perfect inequality (one participant contributing all posts and everyone else contributing none).

6.7.1.3. IS-HANDLED-BY **narrow contributions histogram**
If we plot the activity of each user as a bar plot, sorted left to right by activity, we can assess what proportion of the users are active or not. A narrow peak of high activity implies few people are active.

6.7.2. IS-VIOLATED-BY **inadequate author diversity**
ideas come from “non-disjoint” author sets i.e. where all the authors tend to agree about the pros and cons for different ideas and therefore probably share an intellectual frame

6.7.2.1. IS-HANDLED-BY **ask participants to suggest new members**
Ask contributors to suggest people with alternative views: "fresh blood" for the deliberation.

6.7.2.2. IS-HANDLED-BY **authors have similar rating vectors**
We can perform vector orthogonalization (Householder, 1958) on authors’ rating vectors, followed by a simple vector distance calculation, to assess how much the opinions for different authors diverge.

6.7.2.3. IS-HANDLED-BY **show contributor demographics**
If demographics are available, we can check for diversity, or lack of it, in the participant population.

6.7.2.4. IS-HANDLED-BY **encourage participation from underrepresented demographics**

6.8. HAS-PART **participants have good inputs**
i.e. they are exposed to a diverse range of materials to inform their ideation and decision making

6.8.1. IS-VIOLATED-BY **myopic authoring**
Authors devote themselves to building upon their own contributions without also refining/critiquing content contributed by others.

6.8.1.1. IS-HANDLED-BY **self/other ratio**
measure (N°rating + N° Pro/Con as answer to other's posts)/ N° of Own Post

6.8.2. IS-VIOLATED-BY **miss relevant content**
A user misses content that would elicit more contributions from them, if they had seen it.

6.8.2.1. IS-HANDLED-BY **renewed interest**
Telligent folks (Marc Smith et al) showed that some lurkers (10-20%) would contribute further if they knew that posts that interested them (they had viewed edited rated commented on them) had become “hot” and therefore worth spending time on, so it’s good if a system can notify people when that happens

6.8.2.2. IS-HANDLED-BY **viewed by people w/similar interests**
Use some kind of clustering (e.g. based on vector de-orthogonalization) to find people who have interests like me, and notify me about activity (views, edits, comments, rates, hots) that was interesting to them. do an eigenvector analysis on rating vectors (perhaps we can augment this by taking advantage of tree structure of map?) and look for what was viewed by people located near to you in that eigenspace Maybe one way to look at the problem is that we want to define, for each person, an estimate of the likelihood that they like/dislike each idea. This suggests two thoughts: () people may add both pros and cons for an idea: how do we combine these? Perhaps we need *two* scores for each idea: one p(like), the other p(dislike)? Or we simply say p(like) = 0.5? () if we assign each person a vector that gives the probability that they like each idea, then the default value can be 0.5 (equal chance of liking it or not), so there are no missing values, simplifying the user similarity calculation. We can start with very simple rules for setting the p(like) values, and refine as needed. () if we propagate p(like/dislike) values up the deliberation map, which does seem potentially useful, we may need to use an evidence accumulation math e.g. Bayesian. For example, if I

have two separate lines of evidence for believing conclusion1, the p(conclusion1) is given by the min or max (and not the sum) of P(evidence1) and P(evidence2), depending on whether the evidence pieces have an AND or OR relationship. Currently, the default semantics for all arguments is OR, though I've considered adding AND nodes to allow correct propagation of belief values up the argument map. () how do we propagate p(like) values up past issues? In particular, should we take the issue rating (which is intended to capture the user's estimate of the importance of the issue) into account? I may really like an idea for an unimportant issue, for example: how much impact should that have up the tree?

6.8.2.3. IS-HANDLED-BY **favored post with low/declining ratings**
Lets author know if he/she should try to add arguments that move the community support in the direction I want

6.8.2.4. IS-HANDLED-BY **busy topics**
calculate an activity score, point people towards busy stuff to be relative to other branches, or number of users, or …? (to avoid penning out scores) to be aggregated up tree? Simply keep track of current score plus time since last update – update value whenever a new event occurs, or when read – events propagate up but reduce in strength as they go - make decay rate faster, so gives a current picture of activity

6.8.2.5. IS-HANDLED-BY **near areas that interested me**
Notify user of activity on posts that are "nearby" to posts they have been interested in (viewed, edited, created, rated, hotted) in the past.

6.8.2.6. IS-HANDLED-BY **interesting to people in my social network**
we can notify users of when there is post activity (views, edits, comments, rates, hots) either initiated by, or considered to be interesting to, people in the user's social network

6.8.3. IS-VIOLATED-BY **static subgroups**
(small and) relatively static sets of people work on each part of the deliberation, so there is little "fresh blood", new ideas, new perspectives

6.8.3.1. IS-HANDLED-BY **encourage people to shift topics**
... to break up static groups

6.8.3.2. IS-HANDLED-BY **track subgroups**
... to see whether a small static group has consistently been responsible for all the content in a given topic area in the argument map.

6.8.4. IS-VIOLATED-BY **platform islands**
Community participants use different tools to support online debate and conversations then remain locked within tools. This implies that topics, ideas and outcomes of online conversations remain constrained to specific communities and fail to cross-federate debate across platforms.

6.8.4.1. IS-HANDLED-BY **contributions clustered by platform**
We can compare the argument maps for different platforms to see whether or not contributions are clustered by platform, or not, i.e. whether key content appears in just one or a few of the platforms.

6.8.4.2. IS-HANDLED-BY **"seeders" propagate ideas cross-platform**
"seeders" transfer key ideas from the argument map summary to social media platforms where they did not previously appear.

6.8.5. IS-VIOLATED-BY **balkanization**
“balkanization” means: the community self-organizes into cliques that agree within themselves but disagree with each other. It occurs when a community divides itself into partisan sub-groups where members of each group agree with one other but actively fight against groups with competing ideas. This can be a problem if it means that the sub-groups do not build upon potentially valuable ideas from other groups because of in-group/out-group social dynamics. Cliques form wherein each clique is devoted to a particular class of solutions and either ignores or actively argues against all other ideas, rather than seeing whether new ideas can be created that combine the best features of both.

6.8.5.1. IS-HANDLED-BY **insularity**
There are multiple subgroups (defined by social network analysis) which discuss related topics but do not talk to each other.

6.8.5.2. IS-HANDLED-BY **odd couples**
Bring an author's attention to ideas that come from an author whose interest/rating vector is very *different*, thus fighting balkanization.

6.8.5.3. IS-HANDLED-BY **bias detectors**
Tools are now emerging to detect whether people are one-sided in their news reading (e.g. see http://crowdresearch.org/blog/?p=8244). Perhaps these can be adapted to detect when innovation contributors are one-sided in the inputs they are taking in.

6.8.5.4. IS-HANDLED-BY **attitude space clusters**
for each user, calculate their attitudes towards an idea (i.e. refining idea or uprating or adding pro arguments is positive, downrating or adding con arguments is negative) and (perhaps after singular vector decomposition) look for clusters of distinct groups that are similar within but very different from each other).

6.8.6. IS-ACHIEVED-BY **P: harvest new content**
This is the process wherein harvesters scan through social media in order to find issues, ideas, arguments that can contribute to the social innovation engagement.

6.8.6.1. HAS-PART **P: summarize as map**
Organize content the authors found into a summary argument map.

6.8.6.1.1. HAS-PART **P: unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.8.6.1.1.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.8.6.1.1.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.8.6.1.1.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.8.6.1.2. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.8.6.1.2.1. REQUIRES **find correct tag**

6.8.6.1.2.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.8.6.1.2.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.8.6.1.3. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.8.6.1.3.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.8.6.1.4. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.8.6.1.4.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.8.6.1.4.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.8.6.1.4.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.8.6.1.5. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.8.6.1.5.1. REQUIRES **find correct tag**

6.8.6.1.5.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.8.6.1.5.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.8.6.1.6. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.8.6.1.6.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.8.6.2. REQUIRES **find useful content**
Authors search (e.g. social media) to find (all) content relevant to the social innovation engagement.

6.8.6.2.1. IS-ACHIEVED-BY **sentiment analysis**
use sentiment analysis to troll web to help harvesters find controversy - e.g. negative sentiment probably means a con.

6.8.6.2.2. IS-ACHIEVED-BY **weblinks from fertile sources**
We can mine web link structure to suggest sources e.g. an article section that proposed a pro argument may have links to other pages which are probably cited to support the pro argument

6.8.6.2.3. IS-ACHIEVED-BY **tagged fertile sources**
harvesters can tag social media site pages as fertile or not, so others can also benefit from that resource.

6.8.6.2.4. IS-ACHIEVED-BY **used in argmap**
use argmap post backlinks to suggest fertile social media sources. If a site had good info, go back to see if more is available

6.8.6.2.5. IS-ACHIEVED-BY **find important authors via SNA**
social network analytics (centrality measures and community detection algorithms) can be used to identify highly influential individuals and groups whose outputs may be particularly worthy of harvesting.

6.8.6.2.6. IS-ACHIEVED-BY **topic trends**
Use a trend detection tool such as that provided by google to find hot topics that may be ripe for harvesting.

6.8.6.2.7. IS-ACHIEVED-BY **forwarding statistics**
forwarding relationships in email and microblogs such as twitter can be used to detect possible atoms of interest. Text that is frequently re-forwarded/quoted, for example, might be especially worthy of a harvester's attention

6.8.6.3. HAS-PART **summarize as map**
Organize content the authors found into a summary argument map.

6.8.6.3.1. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.8.6.3.1.1. REQUIRES **avoid duplicates**

Avoid including more than one instance of a given issue, idea, or argument in the map.

6.8.6.3.1.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.8.6.3.1.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.8.6.3.2. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.8.6.3.2.1. REQUIRES **find correct tag**

6.8.6.3.2.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.8.6.3.2.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.8.6.3.3. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.8.6.3.3.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.8.6.3.4. HAS-PART **unbundle into atoms**
unbundle content into argument map atoms tagged with their type: issue, idea, pro, con.

6.8.6.3.4.1. REQUIRES **avoid duplicates**
Avoid including more than one instance of a given issue, idea, or argument in the map.

6.8.6.3.4.1.1. IS-ACHIEVED-BY **document similarity measures**
Document similarity measures (e.g. based on latent semantic analysis) can be used to find redundant atoms of interest, so they can be merged.

6.8.6.3.4.2. REQUIRES **unbundle correctly**
Make sure that content is divided properly into individual argument map elements.

6.8.6.3.5. HAS-PART **tag atoms**
Tag atoms of interest to identify their type (e.g. issue, idea, pro or con, evidence) as well as their topic area.

6.8.6.3.5.1. REQUIRES **find correct tag**

6.8.6.3.5.1.1. IS-ACHIEVED-BY **use reply structures**
structure in harvested social media (e.g. the reply structure in email conversations and threaded web forum and blog comments) can be used to suggest topic tags for atoms. If an atom of interest tagged with topic X came from an email, for example, it suggests that an atom from a reply to that email is likely to address the same topic.

6.8.6.3.5.1.2. IS-ACHIEVED-BY **look for similar tagged atoms**
Document similarity measures and machine learning tools can be used to suggest topic tags for atoms of interest, based on their similarity to previously-tagged atoms.

6.8.6.3.6. HAS-PART **place in map**
This involves adding context i.e. the implicit issue or idea that the given atom logically refers to, in addition to placing the atom in the map.

6.8.6.3.6.1. REQUIRES **place correctly**
Place post in logically correct part of the argument map.

6.9. HAS-PART **participants contribute fully**
participants contribute fully in terms of their time and skills

6.9.1. HAS-PART **critical mass**
There are enough active contributors to create self-sustaining deliberation.

6.9.1.1. IS-ACHIEVED-BY **author/moderator effort ratio**
If moderators account for a disproportionately high fraction of the overall interactions, that suggests that there are not enough participants to achieve a self-sustaining deliberation.

6.9.2. HAS-PART **retain productive contirbutors**
Keep productive contributors involved.

6.9.3. HAS-PART **strong incentives for participation**

6.9.3.1. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.9.3.1.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.9.3.2. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.9.3.2.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.9.3.3. HAS-PART **be a hero**

contributions have a visible impact on a community or problem the user cares about

6.9.3.3.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.9.3.4. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.9.3.4.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.9.3.5. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.9.3.5.1. IS-ACHIEVED-BY **:contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.9.3.6. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.9.3.6.1. IS-ACHIEVED-BY **:use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.9.3.7. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.9.3.7.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.9.3.8. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.9.3.8.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.9.3.9. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.9.3.9.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.9.3.10. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.9.3.10.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.9.3.11. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.9.3.11.1. IS-ACHIEVED-BY **:contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.9.3.12. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.9.3.12.1. IS-ACHIEVED-BY **:use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.9.4. HAS-PART **few disincentives for participation**

6.9.4.1. IS-VIOLATED-BY **abusive behavior**
Abusive behavior by participants can be a disincentive for participation by other participants.

6.9.4.1.1. IS-HANDLED-BY **foul/respectful language**
We can search for language that reveals respectfulness of community interactions e.g. The speaker uses respectful language towards other participants and/or their arguments. Include also moderately respectful language, not only statements such as “your argument is truly brilliant” but also statements such as “your argument is not bad.” The speaker uses foul language to attack other participants on a personal level. Include also mild foul language, not only statements such as “you are a liar” but also statements such as “you seem a little confused.” Code the names of the participants attacked in this way and give the exact quote of the foul language. The speaker uses foul language to attack the arguments of other participants but abstains from personal attacks. Here again include also mild foul language, not only statements such as “this argument is stupid” but also statements such as “this argument is a little weak.” Code the names of the participants whose arguments are attacked in this way and give the exact quote of the foul language.

6.9.5. IS-VIOLATED-BY **one-sided contribution**
participants focus on just a single style of contribution (e.g. adding ideas, or adding arguments) and thus potentially are not contribute some of their skills to the deliberation

6.9.5.1. IS-CAUSED-BY **idea nay-sayer**
User has only critiques, no positive suggestions.

6.9.5.1.1. IS-HANDLED-BY **only cons for issue ideas**
user doesn't like any of existing options for an issue

6.9.5.1.2. IS-HANDLED-BY **suggest alternative**
ask user to suggest a new alternative for an issue

6.9.5.2. IS-HANDLED-BY **six hats stats**
Assess whether the user has used all "six hats":

- Thinking (Blue) - thinking about thinking, process issues
- Information: (White) - considering what information is available, what are the facts? provide details on the issues (= problems to be solved)
- Creativity (Green) - statements of provocation and investigation, seeing where a thought goes (= propose ideas)
- Good points judgment (Yellow) - logic applied to identifying benefits, seeking harmony (= identifying pros)
- Bad points judgment (Black) - logic applied to identifying flaws or barriers, seeking mismatch (= identifying cons)
- Emotions (Red) - instinctive gut reaction or statements of emotional feeling (but not any justification) (=> ratings? comments?) indicative of final stage - detected using prevalence of emotive words?

wikipedia article

6.9.5.3. IS-HANDLED-BY **unjustified contributions**
The participant provides ideas and arguments without providing arguments backing them up. This can have several levels (according to Jurg Steiner's Discourse Quality Index): (1) The speaker does not present any arguments (asks, for example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

6.9.5.4. IS-HANDLED-BY **contributions histogram**
Create histogram assessing user contributions for post types (ideas, issues, arguments) and topics (branches of the map) to detect if they are specializing in a narrow scope.

6.9.6. IS-VIOLATED-BY **non-participation**
Community members are not participating in the deliberation.

6.9.6.1. IS-HANDLED-BY **lapsed contributors**
Detect people who started but stopped participating.

6.9.6.2. IS-HANDLED-BY **user activity stats**

6.9.7. IS-ACHIEVED-BY **calculate contribution scores**
you can add points to capacity building exercises, like if you read x or watch the video or do the training well you get points they assign points to you if you operate in the directions that the organizers wants, so for example when you do a post right you get a point. when you do a mistake you lose one for creation and evaluation and moderation and prediction accuracy You can be offered personalized point-gaining opportunities, identified by the tutoring and attention mediation heuristics i.e. metrics that identify how much a user has contributed. Some possible sources of ideas for how to calculate these scores include: • slashdot (karma points) • digg • yourview.org.au (credibility score) http://www.publicdeliberation.net/jpd/vol8/iss1/art12/ be a thoughtful and constructive participant - particularly in the eyes of other high-credibility participants, and particularly in relation to those who disagree with you. Participate on a wide variety of issues; Participate fully, i.e. by rating, commenting and voting Give due consideration to the relevant arguments on both sides Earn the respect of others, particularly those who already have high credibility, and those who disagree with you. Don't be abusive, rude, obscene, arrogant, or obnoxious.

6.9.7.1. REQUIRES **accurate scores**
The contribution scores should be accurate.

6.9.7.1.1. IS-VIOLATED-BY **reputation gaming**
Users can try to game the system to get high reputation scores without actually contributing much to the social innovation process.

6.9.8. HAS-PART **critical mass**
There are enough active contributors to create self-sustaining deliberation.

6.9.8.1. IS-ACHIEVED-BY **author/moderator effort ratio**
If moderators account for a disproportionately high fraction of the overall interactions, that suggests that there are not enough participants to achieve a self-sustaining deliberation.

6.9.9. HAS-PART **retain productive contirbutors**
Keep productive contributors involved.

6.9.10. HAS-PART **strong incentives for participation**

6.9.10.1. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.9.10.1.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.9.10.2. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.9.10.2.1. IS-ACHIEVED-BY **publicize user contribution scores**

People's contribution scores are generally visible, as for example with stackoverflow points.

6.9.10.3. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.9.10.3.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.9.10.4. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.9.10.4.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.9.10.5. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.9.10.5.1. IS-ACHIEVED-BY **:contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.9.10.6. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.9.10.6.1. IS-ACHIEVED-BY **:use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.9.10.7. HAS-PART **fun**
Users contribute for fun e.g. as a competitive game

6.9.10.7.1. IS-ACHIEVED-BY **leaderboards**
to provide a competitive frame.

6.9.10.8. HAS-PART **reputation**
High-contributing users earn a better reputation, e.g. via badges. This can be translated into real-life benefits (e.g. as in stackoverflow scores helping people get programmer jobs) in addition to just visibility in-system.

6.9.10.8.1. IS-ACHIEVED-BY **publicize user contribution scores**
People's contribution scores are generally visible, as for example with stackoverflow points.

6.9.10.9. HAS-PART **be a hero**
contributions have a visible impact on a community or problem the user cares about

6.9.10.9.1. IS-ACHIEVED-BY **impact metrics**
define metrics that help people see what impact their contributions have had e.g. in terms of eliciting discussion, changing hearts and minds, or even motivating action (if that is tracked).

6.9.10.10. HAS-PART **find your tribe**
Contributing to the system helps you find like-minded people you want to interact with.

6.9.10.10.1. IS-ACHIEVED-BY **tribe-finder metrics**
provide metrics that help people find "members of their trip" based on contributions in the deliberation map

6.9.10.11. HAS-PART **power**
High levels of contributions give the user greater power in the system.

6.9.10.11.1. IS-ACHIEVED-BY **:contribution-based role assignment**
High contribution scores to earn privileges (e.g. allowed to add posts, senior author, moderator, ability to "hot" posts, allowed to edit other's posts) as well as spend points to affect the system e.g. to bring more attention to a topic you care about.

6.9.10.12. HAS-PART **improve system**
The system is able to work better for the user, e.g. because it knows what he/she is interested in, if they contribute - just like with rating movies in Netflix.

6.9.10.12.1. IS-ACHIEVED-BY **:use ratings to filter content**
System filters content to show you just the "good stuff", once it knows what kind of content you prefer. It can do so for looking for content similar/close to content you liked/disliked.

6.9.11. HAS-PART **few disincentives for participation**

6.9.11.1. IS-VIOLATED-BY **abusive behavior**
Abusive behavior by participants can be a disincentive for participation by other participants.

6.9.11.1.1. IS-HANDLED-BY **foul/respectful language**
We can search for language that reveals respectfulness of community interactions e.g. The speaker uses respectful language towards other participants and/or their arguments. Include also moderately respectful language, not only statements such as “your argument is truly brilliant” but also statements such as “your argument is not bad.” The speaker uses foul language to attack other participants on a personal level. Include also mild foul language, not only statements such as “you are a liar” but also statements such as “you seem a little confused.” Code the names of the participants attacked in this way and give the exact quote of the foul language. The speaker uses foul language to attack the arguments of other participants but abstains from personal attacks. Here again include also mild foul language, not only statements such as “this argument is stupid” but also statements such as “this argument is a little weak.” Code the names of the participants whose arguments are attacked in this way and give the exact quote of the foul language.

6.10. HAS-PART **participants contribute effectively**

6.10.1. IS-VIOLATED-BY **troublemakers**
Troublemaker users are reducing the effectiveness of the social innovation system.

6.10.1.1. IS-HANDLED-BY **count trouble tags**
Count trouble tags created in response to actions by that user.

6.10.2. IS-VIOLATED-BY **participants don't know where to contribute**
Participants don't know what parts of the social innovation engagement they can best contribute to.

6.10.2.1. IS-HANDLED-BY **:attention mediation**
The system notifies participants about tasks that need attention and that they are suited to perform.

6.10.2.1.1. HAS-PART **:gather deliberation data**
on both the users and the content they generate.

6.10.2.1.1.1. REQUIRES **sufficient data is available**

6.10.2.1.1.1.1. IS-VIOLATED-BY **insufficient deliberation info**

6.10.2.1.1.1.1.1. IS-HANDLED-BY **:active learning**
Use information theory to determine which actions will have the greater positive impact in terms of identifying the most important exceptions. Thanks to Avi for the idea.

6.10.2.1.2. HAS-PART **:run metrics**
run metrics processes to detect symptoms of exceptions

6.10.2.1.2.1. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.10.2.1.2.2. HAS-PART **:mathematical analysis**

6.10.2.1.2.2.1. HAS-PART **:belief propagation**

6.10.2.1.2.2.2. HAS-PART **:network analysis**

6.10.2.1.2.2.3. HAS-PART **:eigenvector analysis**

6.10.2.1.2.2.4. HAS-PART **:belief propagation**

6.10.2.1.2.2.5. HAS-PART **:network analysis**

6.10.2.1.2.2.6. HAS-PART **:eigenvector analysis**

6.10.2.1.2.3. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.10.2.1.2.4. HAS-PART **:mathematical analysis**

6.10.2.1.2.4.1. HAS-PART **:belief propagation**

6.10.2.1.2.4.2. HAS-PART **:network analysis**

6.10.2.1.2.4.3. HAS-PART **:eigenvector analysis**

6.10.2.1.2.4.4. HAS-PART **:belief propagation**

6.10.2.1.2.4.5. HAS-PART **:network analysis**

6.10.2.1.2.4.6. HAS-PART **:eigenvector analysis**

6.10.2.1.3. HAS-PART **:diagnose exceptions**
Determine which exceptions are taking place, given the current metrics values.

6.10.2.1.4. HAS-PART **:prioritize exceptions**
Since exceptions are identified wrt an "ideal" deliberation process, there are likely to be many possible ones in play , more perhaps than can be handled at one time. For this reason, we need a way to prioritize which exceptions receive more attention.

6.10.2.1.5. HAS-PART **:select handlers**
Select which handler, if there are several available, to enact.

6.10.2.1.5.1. IS-REALIZED-BY **:exception-specific handler**
Pick a handler specific to that exception.

6.10.2.1.5.2. IS-REALIZED-BY **:dynamic incentives**
harvesters can be guided by dynamic pricing whose prizes reflect what customer wants, or what the attention mediation algorithms propose

6.10.2.1.5.3. IS-REALIZED-BY **:notification**
If no more specific handler is available, the default handler is to notify the people who could do something about the exception.

6.10.2.1.6. HAS-PART **:run handlers**
Run the selected handler(s) to address the exception

6.10.2.1.7. HAS-PART **:learn from experience**
i.e. learn how to detect and handle exceptions more effectively over time, based on previous experience.

6.10.2.1.7.1. HAS-PART **:collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.10.2.1.7.2. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.10.2.1.7.3. HAS-PART **:machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.10.2.1.7.4. HAS-PART **:collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.10.2.1.7.5. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.10.2.1.7.6. HAS-PART **:machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.10.2.1.8. HAS-PART **:gather deliberation data**
on both the users and the content they generate.

6.10.2.1.8.1. REQUIRES **sufficient data is available**

6.10.2.1.8.1.1. IS-VIOLATED-BY **insufficient deliberation info**

6.10.2.1.8.1.1.1. IS-HANDLED-BY **:active learning**
Use information theory to determine which actions will have the greater positive impact in terms of identifying the most important exceptions. Thanks to Avi for the idea.

6.10.2.1.9. HAS-PART **:run metrics**
run metrics processes to detect symptoms of exceptions

6.10.2.1.9.1. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.10.2.1.9.2. HAS-PART **:mathematical analysis**

6.10.2.1.9.2.1. HAS-PART **:belief propagation**

6.10.2.1.9.2.2. HAS-PART **:network analysis**

6.10.2.1.9.2.3. HAS-PART **:eigenvector analysis**

6.10.2.1.9.2.4. HAS-PART **:belief propagation**

6.10.2.1.9.2.5. HAS-PART **:network analysis**

6.10.2.1.9.2.6. HAS-PART **:eigenvector analysis**

6.10.2.1.9.3. HAS-PART **:triplestore queries**
i.e. using something like SPARQL or Tinkerpop to run graphical queries over the deliberation data.

6.10.2.1.9.4. HAS-PART **:mathematical analysis**

6.10.2.1.9.4.1. HAS-PART **:belief propagation**

6.10.2.1.9.4.2. HAS-PART **:network analysis**

6.10.2.1.9.4.3. HAS-PART **:eigenvector analysis**

6.10.2.1.9.4.4. HAS-PART **:belief propagation**

6.10.2.1.9.4.5. HAS-PART **:network analysis**

6.10.2.1.9.4.6. HAS-PART **:eigenvector analysis**

6.10.2.1.10. HAS-PART **:diagnose exceptions**
Determine which exceptions are taking place, given the current metrics values.

6.10.2.1.11. HAS-PART **:prioritize exceptions**
Since exceptions are identified wrt an "ideal" deliberation process, there are likely to be many possible ones in play , more perhaps than can be handled at one time. For this reason, we need a way to prioritize which exceptions receive more attention.

6.10.2.1.12. HAS-PART **:select handlers**
Select which handler, if there are several available, to enact.

6.10.2.1.12.1. IS-REALIZED-BY **:exception-specific handler**
Pick a handler specific to that exception.

6.10.2.1.12.2. IS-REALIZED-BY **:dynamic incentives**
harvesters can be guided by dynamic pricing whose prizes reflect what customer wants, or what the attention mediation algorithms propose

6.10.2.1.12.3. IS-REALIZED-BY **:notification**
If no more specific handler is available, the default handler is to notify the people who could do something about the exception.

6.10.2.1.13. HAS-PART **:run handlers**

Run the selected handler(s) to address the exception

6.10.2.1.14. HAS-PART **learn from experience**
i.e. learn how to detect and handle exceptions more effectively over time, based on previous experience.

6.10.2.1.14.1. HAS-PART **collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.10.2.1.14.2. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.10.2.1.14.3. HAS-PART **machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.10.2.1.14.4. HAS-PART **collect user feedback**
Collect feedback, from participants, on the deliberation process e.g. whether or not different interventions (e.g. attention notifications) were helpful or not, whether or not a given discussion is proceeding well, etc. learn best metrics based on relevance feedback from users (i.e. if they thought exceptions were valid, suggestions were useful)

6.10.2.1.14.5. HAS-PART **human visual pattern recognition**
We can create visualizations of deliberations that were tagged, by participants (e.g. as "productive", "contentious", "specialized" etc) and ask humans to figure out which patterns in these visualizations help predict which tags. This can then be the basis of defining new metrics.

6.10.2.1.14.6. HAS-PART **machine learning**
Use machine learning algorithms to find the most important metric and most useful handlers, as well as possibly even learn new better metrics and handlers. use eigenvector analysis to consolidate metrics. • generate new metrics using recombination and mutation, or by (random) walks through graphs, and then finding the ones that best predict our quality measures. feed examples of good and bad argument map fragments to a relational learning algorithm, it can find features that distinguish them - see Abraham Bernstein ESWC-2008 paper on SPARQLML

6.10.3. IS-VIOLATED-BY **users don't understand how to use system**
Users are lacking information or skills needed to use the social innovation system effectively.

6.10.3.1. IS-HANDLED-BY **custom training**
Users can be pointed to help materials that appear relevant to the problems they seem to be having with using the deliberation system correctly. See the reasons that moderators have identified for rejecting posts (e.g. improperly unbundled or located), use this to provide some kind of personalized training for the user - a basis for integrated tutoring system.

6.10.3.2. IS-HANDLED-BY **slow certification**
many moderator iterations are needed for certification of the contributor's posts

6.10.3.3. IS-HANDLED-BY **high action/output ratio**
i.e. users try a lot of edit actions, but produce only a few certified posts as a result

7. REQUIRES **good results**
The social innovation engagement produces good useful results for the customer.

7.1. HAS-PART **complete content**
i.e. the deliberation covers all the content that the customer needs

7.1.1. IS-VIOLATED-BY **incomplete content**
The social innovation missed some important issues, criteria, ideas, or arguments

7.1.1.1. IS-HANDLED-BY **little content for customer "hotted" topics**
Customer indicates that given topics need more attention.

7.1.1.2. IS-HANDLED-BY **little content for peer "hotted" topics**
people add to hot list, score degrades over time and grows with additional endorsements, as a way for contributors to tell peers about what needs attention

7.2. HAS-PART **high-quality content**
The quality of the content is high: important issues, relevant criteria, promising ideas, compelling arguments.

7.2.1. IS-VIOLATED-BY **spam**
Participants have contributed irrelevant material e.g. sales stuff.

7.2.1.1. IS-HANDLED-BY **trouble tags**
Participants can tag posts that they believe has inappropriate, irrelevant, or redundant content.

7.2.1.2. IS-HANDLED-BY **spam detection filters**

7.3. HAS-PART **easy to find what you need**
The desired content of a deliberation can be accessed quickly, easily, and as fully as desired, so it's easy to find the good stuff and know where to contribute new material as well.

7.3.1. IS-VIOLATED-BY **disorganized map**
The summary map for the social innovation is poorly structured, making it hard to find stuff e.g. because it is not arranged as a hierarchical topic tree.

7.3.1.1. IS-HANDLED-BY **measure search times**

Measure how long it takes people to properly place a post whose correct location is already known - if they have to do a lot of searching around to place the post, that suggests the map is poorly organized.

7.3.2. IS-VIOLATED-BY **redundancy**
Where the same or similar ideas are repeated.

7.3.2.1. IS-HANDLED-BY **duplicate posts**
Check for likely duplicates posts in summary map e.g. using LSA or LDA techniques.

7.3.2.2. IS-HANDLED-BY **disincent duplicates**
Reduce reputation/contribution scores for users that contribute duplicates

7.3.3. IS-ACHIEVED-BY **certification process**
This is the process whereby moderators do quality control on the posts contributed by the others to the summary map.

7.3.3.1. HAS-PART **acquire post**
The moderator acquires a post which requires attention - either a pending post or one that has an unaddressed trouble tag.

7.3.3.1.1. REQUIRES **quick certification**
... so contributors are not discouraged by how long it takes for their contributions to get "published"

7.3.3.1.1.1. IS-VIOLATED-BY **long queue times**

7.3.3.1.1.1.1. IS-HANDLED-BY **average wait**
Measure average wait time of pending posts on queue before checked by moderator

7.3.3.2. HAS-PART **check for problems**
The moderator checks post for problems.

7.3.3.2.1. HAS-PART **check location**
Make sure post is in the right part of the summary argument map.

7.3.3.2.2. HAS-PART **check title**
Check if the post title represents a compact accurate summary of post contents

7.3.3.2.3. HAS-PART **check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.3.3.2.4. HAS-PART **check substance**
check if the posts' content is substantive

7.3.3.2.5. HAS-PART **check relevance**
Check if the post is relevant to the topic.

7.3.3.2.6. REQUIRES **complete and accurate**
The moderator finds all the problems, with no false positives.

7.3.3.2.7. HAS-PART **check location**
Make sure post is in the right part of the summary argument map.

7.3.3.2.8. HAS-PART **check title**
Check if the post title represents a compact accurate summary of post contents

7.3.3.2.9. HAS-PART **check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.3.3.2.10. HAS-PART **check substance**
check if the posts' content is substantive

7.3.3.2.11. HAS-PART **check relevance**
Check if the post is relevant to the topic.

7.3.3.3. HAS-PART **take action**
The moderator takes action on the post being checked.

7.3.3.3.1. HAS-PART **fix**
the moderator fixes (some of) the problems in the post

7.3.3.3.2. HAS-PART **certify**
The moderator certifies the post

7.3.3.3.3. HAS-PART **de-certify**
The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.3.3.3.4. HAS-PART **trash**
The moderator trashes the post

7.3.3.3.5. REQUIRES **correct action**
Moderator takes correct action on post.

7.3.3.3.5.1. IS-VIOLATED-BY **moderator bias**
The moderator takes punitive action against a post because the moderator is biased against the content or author of the post.

7.3.3.3.5.1.1. IS-HANDLED-BY **trouble tags**
Participants can post "trouble tags" on posts where they feel moderators are acting improperly.

7.3.3.3.6. HAS-PART **fix**
the moderator fixes (some of) the problems in the post

7.3.3.3.7. HAS-PART **certify**
The moderator certifies the post

7.3.3.3.8. HAS-PART **de-certify**

The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.3.3.3.9. HAS-PART **:trash**
The moderator trashes the post

7.3.3.4. REQUIRES **skilled moderators**
Moderators must be skilled at checking and, if necessary (helping authors to) fix posts so they can be certified.

7.3.3.5. HAS-PART **:acquire post**
The moderator acquires a post which requires attention - either a pending post or one that has an unaddressed trouble tag.

7.3.3.5.1. REQUIRES **quick certification**
... so contributors are not discouraged by how long it takes for their contributions to get "published"

7.3.3.5.1.1. IS-VIOLATED-BY **long queue times**

7.3.3.5.1.1.1. IS-HANDLED-BY **average wait**
Measure average wait time of pending posts on queue before checked by moderator

7.3.3.6. HAS-PART **:check for problems**
The moderator checks post for problems.

7.3.3.6.1. HAS-PART **:check location**
Make sure post is in the right part of the summary argument map.

7.3.3.6.2. HAS-PART **:check title**
Check if the post title represents a compact accurate summary of post contents

7.3.3.6.3. HAS-PART **:check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.3.3.6.4. HAS-PART **:check substance**
check if the posts' content is substantive

7.3.3.6.5. HAS-PART **:check relevance**
Check if the post is relevant to the topic.

7.3.3.6.6. REQUIRES **complete and accurate**
The moderator finds all the problems, with no false positives.

7.3.3.6.7. HAS-PART **:check location**
Make sure post is in the right part of the summary argument map.

7.3.3.6.8. HAS-PART **:check title**
Check if the post title represents a compact accurate summary of post contents

7.3.3.6.9. HAS-PART **:check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.3.3.6.10. HAS-PART **:check substance**
check if the posts' content is substantive

7.3.3.6.11. HAS-PART **:check relevance**
Check if the post is relevant to the topic.

7.3.3.7. HAS-PART **:take action**
The moderator takes action on the post being checked.

7.3.3.7.1. HAS-PART **:fix**
the moderator fixes (some of) the problems in the post

7.3.3.7.2. HAS-PART **:certify**
The moderator certifies the post

7.3.3.7.3. HAS-PART **:de-certify**
The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.3.3.7.4. HAS-PART **:trash**
The moderator trashes the post

7.3.3.7.5. REQUIRES **correct action**
Moderator takes correct action on post.

7.3.3.7.5.1. IS-VIOLATED-BY **moderator bias**
The moderator takes punitive action against a post because the moderator is biased against the content or author of the post.

7.3.3.7.5.1.1. IS-HANDLED-BY **trouble tags**
Participants can post "trouble tags" on posts where they feel moderators are acting improperly.

7.3.3.7.6. HAS-PART **:fix**
the moderator fixes (some of) the problems in the post

7.3.3.7.7. HAS-PART **:certify**
The moderator certifies the post

7.3.3.7.8. HAS-PART **:de-certify**
The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.3.3.7.9. HAS-PART **:trash**
The moderator trashes the post

7.4. HAS-PART **complete content**
i.e. the deliberation covers all the content that the customer needs

7.4.1. IS-VIOLATED-BY **incomplete content**
The social innovation missed some important issues, criteria, ideas, or arguments

7.4.1.1. IS-HANDLED-BY **little content for customer "hotted" topics**
Customer indicates that given topics need more attention.

7.4.1.2. IS-HANDLED-BY **little content for peer "hotted" topics**
people add to hot list, score degrades over time and grows with additional endorsements, as a way for contributors to tell peers about what needs attention

7.5. HAS-PART **high-quality content**
The quality of the content is high: important issues, relevant criteria, promising ideas, compelling arguments.

7.5.1. IS-VIOLATED-BY **spam**
Participants have contributed irrelevant material e.g. sales stuff.

7.5.1.1. IS-HANDLED-BY **trouble tags**
Participants can tag posts that they believe has inappropriate, irrelevant, or redundant content.

7.5.1.2. IS-HANDLED-BY **spam detection filters**

7.6. HAS-PART **easy to find what you need**
The desired content of a deliberation can be accessed quickly, easily, and as fully as desired, so it's easy to find the good stuff and know where to contribute new material as well.

7.6.1. IS-VIOLATED-BY **disorganized map**
The summary map for the social innovation is poorly structured, making it hard to find stuff e.g. because it is not arranged as a hierarchical topic tree.

7.6.1.1. IS-HANDLED-BY **measure search times**
Measure how long it takes people to properly place a post whose correct location is already known - if they have to do a lot of searching around to place the post, that suggests the map is poorly organized.

7.6.2. IS-VIOLATED-BY **redundancy**
Where the same or similar ideas are repeated.

7.6.2.1. IS-HANDLED-BY **duplicate posts**
Check for likely duplicates posts in summary map e.g. using LSA or LDA techniques.

7.6.2.2. IS-HANDLED-BY **disincent duplicates**
Reduce reputation/contribution scores for users that contribute duplicates

7.6.3. IS-ACHIEVED-BY **certification process**
This is the process whereby moderators do quality control on the posts contributed by the others to the summary map.

7.6.3.1. HAS-PART **acquire post**
The moderator acquires a post which requires attention - either a pending post or one that has an unaddressed trouble tag.

7.6.3.1.1. REQUIRES **quick certification**
... so contributors are not discouraged by how long it takes for their contributions to get "published"

7.6.3.1.1.1. IS-VIOLATED-BY **long queue times**

7.6.3.1.1.1.1. IS-HANDLED-BY **average wait**
Measure average wait time of pending posts on queue before checked by moderator

7.6.3.2. HAS-PART **check for problems**
The moderator checks post for problems.

7.6.3.2.1. HAS-PART **check location**
Make sure post is in the right part of the summary argument map.

7.6.3.2.2. HAS-PART **check title**
Check if the post title represents a compact accurate summary of post contents

7.6.3.2.3. HAS-PART **check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.6.3.2.4. HAS-PART **check substance**
check if the posts' content is substantive

7.6.3.2.5. HAS-PART **check relevance**
Check if the post is relevant to the topic.

7.6.3.2.6. REQUIRES **complete and accurate**
The moderator finds all the problems, with no false positives.

7.6.3.2.7. HAS-PART **check location**
Make sure post is in the right part of the summary argument map.

7.6.3.2.8. HAS-PART **check title**
Check if the post title represents a compact accurate summary of post contents

7.6.3.2.9. HAS-PART **check unbundling**
Check if the post represents a single issue, criterion, idea, or argument.

7.6.3.2.10. HAS-PART **check substance**
check if the posts' content is substantive

7.6.3.2.11. HAS-PART **check relevance**
Check if the post is relevant to the topic.

7.6.3.3. HAS-PART **take action**
The moderator takes action on the post being checked.

- 7.6.3.3.1. HAS-PART **fix**
  the moderator fixes (some of) the problems in the post
- 7.6.3.3.2. HAS-PART **certify**
  The moderator certifies the post
- 7.6.3.3.3. HAS-PART **de-certify**
  The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).
- 7.6.3.3.4. HAS-PART **trash**
  The moderator trashes the post
- 7.6.3.3.5. REQUIRES **correct action**
  Moderator takes correct action on post.
  - 7.6.3.3.5.1. IS-VIOLATED-BY **moderator bias**
    The moderator takes punitive action against a post because the moderator is biased against the content or author of the post.
    - 7.6.3.3.5.1.1. IS-HANDLED-BY **trouble tags**
      Participants can post "trouble tags" on posts where they feel moderators are acting improperly.
- 7.6.3.3.6. HAS-PART **fix**
  the moderator fixes (some of) the problems in the post
- 7.6.3.3.7. HAS-PART **certify**
  The moderator certifies the post
- 7.6.3.3.8. HAS-PART **de-certify**
  The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).
- 7.6.3.3.9. HAS-PART **trash**
  The moderator trashes the post

7.6.3.4. REQUIRES **skilled moderators**
Moderators must be skilled at checking and, if necessary (helping authors to) fix posts so they can be certified.

7.6.3.5. HAS-PART **acquire post**
The moderator acquires a post which requires attention - either a pending post or one that has an unaddressed trouble tag.

- 7.6.3.5.1. REQUIRES **quick certification**
  ... so contributors are not discouraged by how long it takes for their contributions to get "published"
  - 7.6.3.5.1.1. IS-VIOLATED-BY **long queue times**
    - 7.6.3.5.1.1.1. IS-HANDLED-BY **average wait**
      Measure average wait time of pending posts on queue before checked by moderator

7.6.3.6. HAS-PART **check for problems**
The moderator checks post for problems.

- 7.6.3.6.1. HAS-PART **check location**
  Make sure post is in the right part of the summary argument map.
- 7.6.3.6.2. HAS-PART **check title**
  Check if the post title represents a compact accurate summary of post contents
- 7.6.3.6.3. HAS-PART **check unbundling**
  Check if the post represents a single issue, criterion, idea, or argument.
- 7.6.3.6.4. HAS-PART **check substance**
  check if the posts' content is substantive
- 7.6.3.6.5. HAS-PART **check relevance**
  Check if the post is relevant to the topic.
- 7.6.3.6.6. REQUIRES **complete and accurate**
  The moderator finds all the problems, with no false positives.
- 7.6.3.6.7. HAS-PART **check location**
  Make sure post is in the right part of the summary argument map.
- 7.6.3.6.8. HAS-PART **check title**
  Check if the post title represents a compact accurate summary of post contents
- 7.6.3.6.9. HAS-PART **check unbundling**
  Check if the post represents a single issue, criterion, idea, or argument.
- 7.6.3.6.10. HAS-PART **check substance**
  check if the posts' content is substantive
- 7.6.3.6.11. HAS-PART **check relevance**
  Check if the post is relevant to the topic.

7.6.3.7. HAS-PART **take action**
The moderator takes action on the post being checked.

- 7.6.3.7.1. HAS-PART **fix**
  the moderator fixes (some of) the problems in the post

7.6.3.7.2. HAS-PART **certify**
The moderator certifies the post

7.6.3.7.3. HAS-PART **de-certify**
The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.6.3.7.4. HAS-PART **trash**
The moderator trashes the post

7.6.3.7.5. REQUIRES **correct action**
Moderator takes correct action on post.

7.6.3.7.5.1. IS-VIOLATED-BY **moderator bias**
The moderator takes punitive action against a post because the moderator is biased against the content or author of the post.

7.6.3.7.5.1.1. IS-HANDLED-BY **trouble tags**
Participants can post "trouble tags" on posts where they feel moderators are acting improperly.

7.6.3.7.6. HAS-PART **fix**
the moderator fixes (some of) the problems in the post

7.6.3.7.7. HAS-PART **certify**
The moderator certifies the post

7.6.3.7.8. HAS-PART **de-certify**
The moderator de-certifies the post because it no longer meets the requirements (but leaves it in the map so an author can fix it).

7.6.3.7.9. HAS-PART **trash**
The moderator trashes the post

8. REQUIRES **side benefits**
The social innovation system has positive impacts other than simply the content it generates e.g. in terms of impact on the participants.

8.1. HAS-PART **strengthen community**
The user community becomes stronger, i.e. better able to deal with future challenges.

8.1.1. HAS-PART **skill development**
The community, by virtue of participating in the social innovation system, develops skills that will be valuable in the future.

8.1.2. HAS-PART **increased connection**
The community members develop an increased sense of connection and trust in each other, which can be helpful for future shared tasks.

8.1.2.1. IS-VIOLATED-BY **upward acrimony spiral**
The level of contention in the community grows due to conflict among members.

8.1.2.1.1. IS-HANDLED-BY **longitudinal sentiment analysis**
use longitudinal sentiment analysis to detect the trends in "hot talk" e.g. in the form of strong adjectives and swear words and antonyms see: Determining Modality and Factuality for Text Entailment

8.1.2.2. IS-VIOLATED-BY **divergent vocabularies**
The vocabulary used by the participants diverges over time, suggesting that they are not reaching the common ground needed for effective collaboation.

8.1.2.2.1. IS-HANDLED-BY **common ground vocabulary**
assess trends in how much vocabulary is shared across authors

8.1.3. HAS-PART **greater consensus**
The community develops a higher level of consensus in terms of issues, criteria, and so on.

8.1.3.1. IS-VIOLATED-BY **polarization**
Participant attitudes toward ideas are becoming more, rather than less, divergent.

8.1.3.1.1. IS-HANDLED-BY **bimodal ratings histogram**
This can be calculated in different ways, e.g. o ratings variance across competing ideas increases over time o average controversy scores grow over time

8.1.4. HAS-PART **improved connectivity**
The community members learn more about who the other members are, and what they can do, so they know who to call on for what purpose in the future.

8.1.4.1. IS-VIOLATED-BY **un-small worlds**
It is difficult for participants to find people with relevant interests

8.1.4.1.1. IS-HANDLED-BY **social network analysis**
look at interaction topology, e.g. for clumps of poorly/non-connected users

8.1.4.1.2. IS-HANDLED-BY **user surveys**

8.1.5. HAS-PART **skill development**
The community, by virtue of participating in the social innovation system, develops skills that will be valuable in the future.

8.1.6. HAS-PART **increased connection**
The community members develop an increased sense of connection and trust in each other, which can be helpful for future shared tasks.

8.1.6.1. IS-VIOLATED-BY **upward acrimony spiral**
The level of contention in the community grows due to conflict among members.

8.1.6.1.1. IS-HANDLED-BY **longitudinal sentiment analysis**

use longitudinal sentiment analysis to detect the trends in "hot talk" e.g. in the form of strong adjectives and swear words and antonyms see: Determining Modality and Factuality for Text Entailment

8.1.6.2. IS-VIOLATED-BY **divergent vocabularies**
The vocabulary used by the participants diverges over time, suggesting that they are not reaching the common ground needed for effective collaboation.

8.1.6.2.1. IS-HANDLED-BY **common ground vocabulary**
assess trends in how much vocabulary is shared across authors

8.1.7. HAS-PART **greater consensus**
The community develops a higher level of consensus in terms of issues, criteria, and so on.

8.1.7.1. IS-VIOLATED-BY **polarization**
Participant attitudes toward ideas are becoming more, rather than less, divergent.

8.1.7.1.1. IS-HANDLED-BY **bimodal ratings histogram**
This can be calculated in different ways, e.g. o ratings variance across competing ideas increases over time o average controversy scores grow over time

8.1.8. HAS-PART **improved connectivity**
The community members learn more about who the other members are, and what they can do, so they know who to call on for what purpose in the future.

8.1.8.1. IS-VIOLATED-BY **un-small worlds**
It is difficult for participants to find people with relevant interests

8.1.8.1.1. IS-HANDLED-BY **social network analysis**
look at interaction topology, e.g. for clumps of poorly/non-connected users

8.1.8.1.2. IS-HANDLED-BY **user surveys**

8.2. HAS-PART **learn about community**
One goal of running deliberations is to get an idea of which people in a community are like in terms of their skills and styles, for future reference. We might use this information, for example, to set up filters to visualize and search the content maps (e.g. show me only what 4-stars experts think or discard liberal people, etc.)

8.2.1. IS-ACHIEVED-BY **self-tagging**
We can set up user profiles. In part such profiles could be based on voluntarily provided information upon registration like in facebook (e.g. through variables like, sex, age, religion, political orientation, profession, etc. – see the "Who I am" example in IBM beehive). We could think this voluntary profiling as a kind of self-social tagging.

8.2.2. IS-ACHIEVED-BY **super-users**
Find super-users - this with a top-decile level of activity.

8.2.3. IS-ACHIEVED-BY **who likes what**
We can learn who supports/attacks what. More often than not people value information depending on how much they are able to recognize and trust the source.

8.2.4. IS-ACHIEVED-BY **who knows who**
An additional source of information for profiling can come from the analysis of the social network to look for things like central users.

8.2.5. IS-ACHIEVED-BY **who knows what**
The system can produce a "who knows what map", which could be used for example to find the experts and as an incentives for authors (who are the 4 stars experts in solar energy in this community?). This can be based on authorship info, post ratings, and content classification. we can do social network analysis to see which users are central in which topics. We can assess how widely read a person's contributions are, e.g. in terms of # views #ratings #edits #comments We can look for authors whose content has survived over the course of multiple edits to posts - cf http//trust.cse.ucsc.edu/.

8.3. HAS-PART **strengthen community**
The user community becomes stronger, i.e. better able to deal with future challenges.

8.3.1. HAS-PART **skill development**
The community, by virtue of participating in the social innovation system, develops skills that will be valuable in the future.

8.3.2. HAS-PART **increased connection**
The community members develop an increased sense of connection and trust in each other, which can be helpful for future shared tasks.

8.3.2.1. IS-VIOLATED-BY **upward acrimony spiral**
The level of contention in the community grows due to conflict among members.

8.3.2.1.1. IS-HANDLED-BY **longitudinal sentiment analysis**
use longitudinal sentiment analysis to detect the trends in "hot talk" e.g. in the form of strong adjectives and swear words and antonyms see: Determining Modality and Factuality for Text Entailment

8.3.2.2. IS-VIOLATED-BY **divergent vocabularies**
The vocabulary used by the participants diverges over time, suggesting that they are not reaching the common ground needed for effective collaboation.

8.3.2.2.1. IS-HANDLED-BY **common ground vocabulary**
assess trends in how much vocabulary is shared across authors

8.3.3. HAS-PART **greater consensus**
The community develops a higher level of consensus in terms of issues, criteria, and so on.

8.3.3.1. IS-VIOLATED-BY **polarization**
Participant attitudes toward ideas are becoming more, rather than less, divergent.

8.3.3.1.1. IS-HANDLED-BY **bimodal ratings histogram**

This can be calculated in different ways, e.g. o ratings variance across competing ideas increases over time o average controversy scores grow over time

8.3.4. HAS-PART **improved connectivity**
The community members learn more about who the other members are, and what they can do, so they know who to call on for what purpose in the future.

8.3.4.1. IS-VIOLATED-BY **un-small worlds**
It is difficult for participants to find people with relevant interests

8.3.4.1.1. IS-HANDLED-BY **social network analysis**
look at interaction topology, e.g. for clumps of poorly/non-connected users

8.3.4.1.2. IS-HANDLED-BY **user surveys**

8.3.5. HAS-PART **skill development**
The community, by virtue of participating in the social innovation system, develops skills that will be valuable in the future.

8.3.6. HAS-PART **increased connection**
The community members develop an increased sense of connection and trust in each other, which can be helpful for future shared tasks.

8.3.6.1. IS-VIOLATED-BY **upward acrimony spiral**
The level of contention in the community grows due to conflict among members.

8.3.6.1.1. IS-HANDLED-BY **longitudinal sentiment analysis**
use longitudinal sentiment analysis to detect the trends in "hot talk" e.g. in the form of strong adjectives and swear words and antonyms see: Determining Modality and Factuality for Text Entailment

8.3.6.2. IS-VIOLATED-BY **divergent vocabularies**
The vocabulary used by the participants diverges over time, suggesting that they are not reaching the common ground needed for effective collaboation.

8.3.6.2.1. IS-HANDLED-BY **common ground vocabulary**
assess trends in how much vocabulary is shared across authors

8.3.7. HAS-PART **greater consensus**
The community develops a higher level of consensus in terms of issues, criteria, and so on.

8.3.7.1. IS-VIOLATED-BY **polarization**
Participant attitudes toward ideas are becoming more, rather than less, divergent.

8.3.7.1.1. IS-HANDLED-BY **bimodal ratings histogram**
This can be calculated in different ways, e.g. o ratings variance across competing ideas increases over time o average controversy scores grow over time

8.3.8. HAS-PART **improved connectivity**
The community members learn more about who the other members are, and what they can do, so they know who to call on for what purpose in the future.

8.3.8.1. IS-VIOLATED-BY **un-small worlds**
It is difficult for participants to find people with relevant interests

8.3.8.1.1. IS-HANDLED-BY **social network analysis**
look at interaction topology, e.g. for clumps of poorly/non-connected users

8.3.8.1.2. IS-HANDLED-BY **user surveys**

8.4. HAS-PART **learn about community**
One goal of running deliberations is to get an idea of which people in a community are like in terms of their skills and styles, for future reference. We might use this information, for example, to set up filters to visualize and search the content maps (e.g. show me only what 4-stars experts think or discard liberal people, etc.)

8.4.1. IS-ACHIEVED-BY **self-tagging**
We can set up user profiles. In part such profiles could be based on voluntarily provided information upon registration like in facebook (e.g. through variables like, sex, age, religion, political orientation, profession, etc. – see the “Who I am” example in IBM beehive). We could think this voluntary profiling as a kind of self-social tagging.

8.4.2. IS-ACHIEVED-BY **super-users**
Find super-users - this with a top-decile level of activity.

8.4.3. IS-ACHIEVED-BY **who likes what**
We can learn who supports/attacks what. More often than not people value information depending on how much they are able to recognize and trust the source.

8.4.4. IS-ACHIEVED-BY **who knows who**
An additional source of information for profiling can come from the analysis of the social network to look for things like central users.

8.4.5. IS-ACHIEVED-BY **who knows what**
The system can produce a “who knows what map”, which could be used for example to find the experts and as an incentives for authors (who are the 4 stars experts in solar energy in this community?). This can be based on authorship info, post ratings, and content classification. we can do social network analysis to see which users are central in which topics. We can assess how widely read a person’s contributions are, e.g. in terms of # views #ratings #edits #comments We can look for authors whose content has survived over the course of multiple edits to posts - cf http//trust.cse.ucsc.edu/.

9. HAS-PART **define problem(s)**
Define the problem(s) that the social innovation engagement is supposed to solve.

9.1. REQUIRES **identify issues**
Describe the issues that need to be solved e.g. "what can we do to solve climate change"?

9.1.1. IS-VIOLATED-BY **missing key issues**

9.1.1.1. IS-HANDLED-BY **ask experts**
ask experts to pre-populate map with all key issues

9.1.1.2. IS-HANDLED-BY **expert evaluation**
Experts assess whether or not map includes all key issues.

9.2. REQUIRES **identify criteria**
identify the attributes of a good solution to the problem e.g. "limit average global temperature rise to no more than 2 degrees celsius "

9.2.1. HAS-PART **identify *only* relevant criteria**
Identify only criteria that are relevant/important for this problem - no false positives.

9.2.1.1. IS-VIOLATED-BY **irrelevant criteria**

9.2.1.1.1. IS-HANDLED-BY **low rating**
Criterion has a low rating score.

9.2.2. HAS-PART **identify all criteria**
Identify all relevant/important criteria for this problem.

9.2.2.1. IS-VIOLATED-BY **parochial criteria**
The participants identify only criteria that impact their own self-interest and not common good criteria such as social welfare and equality.

9.2.3. HAS-PART **identify *only* relevant criteria**
Identify only criteria that are relevant/important for this problem - no false positives.

9.2.3.1. IS-VIOLATED-BY **irrelevant criteria**

9.2.3.1.1. IS-HANDLED-BY **low rating**
Criterion has a low rating score.

9.2.4. HAS-PART **identify all criteria**
Identify all relevant/important criteria for this problem.

9.2.4.1. IS-VIOLATED-BY **parochial criteria**
The participants identify only criteria that impact their own self-interest and not common good criteria such as social welfare and equality.

10. HAS-PART **identify solutions**
Identify candidate solutions for the identified problems.

10.1. REQUIRES **high-quality ideas**
Existing social media tend to elicit lots of shallow ideas, with highly variable quality and originality. How can we maximize the proportion of creative, high-quality, deeply considered ideas?

10.1.1. IS-VIOLATED-BY **idea sabotage**
People who don't like an idea edit it to make it worse.

10.1.1.1. IS-HANDLED-BY **edit wars by partisans**
i.e. where someone who doesn't like idea is editing it in conflict with someone who likes the idea. In other words, look for alternating edits by users that appear to have divergent opinions (based on their rating behavior) about the issue they are proposing solutions for.

10.1.1.2. IS-HANDLED-BY **hub-and-spoke interaction network**
Secretive sabotage communication patterns tend towards a hub-and-spoke architecture, as opposed to the network-topology connectivity that characterizes full open discussion. See: Brandy Aven (2011). The effect of corruption on organizational networks and individual behavior. Proceedings of the MIT WIDS colloquium (http//wids.lids.mit.edu/).

10.1.2. IS-VIOLATED-BY **solo ideation**
Authors do not collaborate to refine ideas.

10.1.2.1. IS-HANDLED-BY **single editor**
This post has had only one editor (not counting moderators).

10.1.3. IS-VIOLATED-BY **insular ideation**
ideas do not build upon one another

10.1.3.1. IS-HANDLED-BY **no common ground vocabulary**
look for growing use of shared words and word clusters within topics, which is a way of assessing whether people are building ideas by re-combining existing ones.

10.1.4. IS-ACHIEVED-BY **low idea ratings**
The ideas receive low average "promising" rating from the community.

10.2. REQUIRES **complete idea space**
We want to have a comprehensive picture of the most promising solutions for the problems focused on by the innovation engagement.

10.2.1. IS-VIOLATED-BY **incomplete idea coverage**
The deliberation has only incompletely covered the space of potentially relevant ideas for an (important) issue.

10.2.1.1. IS-HANDLED-BY **attention/importance ratio**
Measure the ratio of attention to issue importance for issues, and highlight the issues with particularly low scores. Issue importance can be calculated by accounting for the importance of the parent issues and promise of the parent ideas.

10.2.1.2. IS-HANDLED-BY **get n(idea) estimates**
Ask users to estimate how many good ideas there are for each issue (e.g. whenever someone creates an issue, or adds an idea to an issue, or even views an issue). The average of that number gives the standard, and we flag an exception if we are substantially below that number of ideas for an issue.

10.2.1.3. IS-HANDLED-BY **ask expert panel**

An expert panel assess whether or not the idea space is covered fully (for a given issue).

10.2.2. IS-VIOLATED-BY **creativity stagnation**
Few novel/interesting ideas are being generated as proposed solutions for a problem.

10.2.2.1. IS-CAUSED-BY **idea groupthink**
groupthink can be defined as a group dedicating the bulk of its attention to refining a single idea, often the first one endorsed by an influential figure, rather than comparing several alternatives in depth.

10.2.2.1.1. IS-HANDLED-BY **attention narrowing**
we can measure when one idea under an issue is receiving the bulk of the community's attention (views, rates, edits and additions) while competing ideas and their underlying arguments remain largely untouched

10.2.2.2. IS-HANDLED-BY **count # ideas rated as novel**
Participants assign a degree of novelty (High, Average and Low) to the posted ideas. The degree of novelty of an issue is the max value attributed to each idea associated to that issue. The degree of novelty of an idea is the average degree of novelty assigned by the crowd to the idea.

10.2.2.3. IS-HANDLED-BY **low vocabulary diversity**
We can measure the use of shared vocabulary in the ideas for a given issue. If there is heavy use of shared terminology, this suggests that the ideas are only moderately diverse. Ideas that are truly diverse will tend to use different vocabulary to express them. In other words, look for ideas that are quite different, in terms of the word frequency statistics, from the other ideas (for that part of the map). We can use LSA or LDA or other document similarity algorithms for this purpose.

10.2.2.4. IS-HANDLED-BY **idea GA**
Usually a solution consists of a *package* of interrelated ideas, so the complete solution space will consist of different combinations of "atomic" ideas. These recombinant space can of course be vast, however, so in practice we must focus on only "promising" packages if at all possible. We can use a GA approach to draw people towards posts. People can score posts on creativity vs. practicality, and weight creativity more at first, practicality more as we near the final point using latent semantic analysis to help identify out-of-box posts and give them higher fitness scores - to maintain diversity. The system can also point people to pairs of ideas - e.g. ideas for different parts of a system, or different ideas for the same subsystem – and suggest they create a new idea based on these existing ones. This has the advantage that we interleave generation and evaluation to help produce a more efficient process (as opposed to generate everything first, and then evaluate the whole redundant mess). The system can suggest users look at combinations that will speed that search for optimal idea combinations when issues are interdependent and utility functions are therefore nonlinear. This can be based on techniques for simulated annealing, creating sub-negotiations for tightly-interdependent issue clusters, etc.

10.2.2.5. IS-HANDLED-BY **"red herrings"**
Use "out of the box" prompts to help break a creative deadlock e.g.: • oblique strategy cards (phrases or cryptic remarks) • randomly selected ideas from the summary map • ideas selected from areas/people the author has heretofore ignored

11. HAS-PART **evaluate solutions**
evaluate solutions with respect to the goals identified for the deliberation

11.1. REQUIRES **high-quality evaluation**
The evaluation provides accurate assessments of the worth of proposed solutions.

11.1.1. HAS-PART **users understand content**
Users understand the content of the map well enough to offer informed evaluations of the ideas described therein.

11.1.1.1. IS-ACHIEVED-BY **narrative summaries**
Convert argmaps into easy-to-follow narrative summaries that make it easier for evaluators to see the key points they need to. Perhaps we can do so taking advantage of rhetorical structure theory?

11.1.2. HAS-PART **complete argumentation**
i.e. the evaluation includes a comprehensive overview of the arguments for and against each proposed solution idea

11.1.2.1. IS-VIOLATED-BY **missing arguments**
An idea is missing some important arguments for or against it.

11.1.2.1.1. IS-CAUSED-BY **self-focused**
The participants generate arguments that refer to criteria that concern them personally, but not those that impact welfare for other groups.

11.1.2.1.2. IS-HANDLED-BY **neglected criteria**
Few/no arguments have been created that evaluate a given idea with respect to one of the solution criteria for that problem.

11.1.2.1.3. IS-HANDLED-BY **few/no arguments**
i.e. there are (important) ideas that have few or no arguments attached to them

11.1.2.1.4. IS-HANDLED-BY **unbalanced arguments**
There is a large imbalance in the number of pros and cons in the debate over an idea.

11.1.2.1.5. IS-HANDLED-BY **few people contributed arguments**

11.1.2.1.6. IS-HANDLED-BY **idea/argument rating disconnect**
We can use such techniques as Bayesian inference (Bolstad, 2010) to propagate a user's ratings for arguments up the argument map to predict how the user should have rated the ideas these arguments address. If there is a large divergence between a user's predicted and actual ratings for an idea, that suggests that the user has not yet entered arguments that are compelling to him or her. A suggestion is to show users the possible misalignment between popularity ratings assigned freely by users and computed scores based on the structure of arguments and how much underlying pros and cons are supported. The gap can be used also as an incentive to users: a popular idea with poor computed score should invite its supporters to provide new arguments or improve the existing ones.

11.1.3. HAS-PART **high-quality argumentation**
The arguments entered in the summary map are well-founded.

11.1.3.1. IS-VIOLATED-BY **false premises**

the arguments made are based on false premises

11.1.3.1.1. IS-CAUSED-BY **argument sabotage**
Someone who disagrees with an argument edits it to sabotage it, rather than (properly) simply downrating or arguing against it.

11.1.3.1.1.1. IS-HANDLED-BY **argument edit wars**
assess the prevalence of edit wars (rapid alternating rollbacks) in the post edit histories – esp. by people who take differing positions on the issues they are editing arguments for. see ZIF workshop paper by János Kertész, Budapest: Edit wars on the Wikipedia:. an interesting measure for wikipedia article controversiality = # mutual reverts by more established user accounts (discounting young vandals)

11.1.3.1.2. IS-HANDLED-BY **low argument ratings**
An argument got a low average rating from the community.

11.1.3.1.3. IS-HANDLED-BY **expert evaluation**
Expert(s) judged that the argument is ill-founded.

11.1.3.2. IS-VIOLATED-BY **incorrect inference**
the arguments made are based on logical fallacies

11.1.3.2.1. IS-HANDLED-BY **automated feedback**
The use of artificial intelligence techniques holds promise to increase the effectiveness of argumentation systems by automatically analyzing user actions and providing supportive feedback. see: http//www.ascilite.org.au/ajet/ajet25/butchart.pdf Mclaren, B. M., Scheuer, O., & Mikšátko, J. (2010). Supporting collaborative learning and e-Discussions using artificial intelligence techniques. International Journal of Artificial Intelligence in Education, 20(1)(1), 1-46. Scheuer, O., McLaren, B. M., Loll, F., & Pinkwart, N. (2012). Automated Analysis and Feedback Techniques to Support Argumentation: A Survey. In: N. Pinkwart, & B. M. McLaren (Eds.), Educational Technologies for Teaching Argumentation Skills (pp. 71–124). Bentham Science Publishers.

11.1.3.2.2. IS-HANDLED-BY **expert judgment**
Expert(s) judge that the inference is faulty.

11.1.4. HAS-PART **high-quality ratings**
The community ratings for the argument map contents (i.e. whether issues are relevant, criteria are important, ideas are promising, arguments are compelling) are accurate.

11.1.4.1. IS-VIOLATED-BY **too few ratings**
There are too few ratings for a post to draw reliable conclusions about how the community judges it's value.

11.1.4.2. IS-VIOLATED-BY **dishonest ratings**
Ratings are dishonest.

11.1.4.2.1. IS-HANDLED-BY **rating inconsistency**
A participant gives inconsistent ratings e.g. they rate the arguments supporting an idea highly, but give the idea itself a poor rating.

11.1.4.3. IS-VIOLATED-BY **incorrect ratings**
the user's ratings for the posts are incorrect

11.1.4.3.1. IS-HANDLED-BY **missing suppoprt**
someone gave a strong + or - rating without a backup argument that they authored or highly rated. This can take several levels, according to the Discourse Quality Index (Jurgen Steiner): (1) The speaker does not present any arguments (asks, for example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

11.1.4.3.2. IS-HANDLED-BY **ignored arguments**
User did not attend to (read, rate) arguments when rating posts impacted by these arguments

11.1.4.3.3. IS-HANDLED-BY **irrational ratings**
see how well a model of rational rating predicts user's ratings. We can use knowledge about map structure to infer what a user's rating for a post "should" be given their ratings for posts below it in tree. For example, with mutex ideas, a rate for X is probably a rate against alternative to X.

11.1.5. HAS-PART **users understand content**
Users understand the content of the map well enough to offer informed evaluations of the ideas described therein.

11.1.5.1. IS-ACHIEVED-BY **narrative summaries**
Convert argmaps into easy-to-follow narrative summaries that make it easier for evaluators to see the key points they need to. Perhaps we can do so taking advantage of rhetorical structure theory?

11.1.6. HAS-PART **complete argumentation**
i.e. the evaluation includes a comprehensive overview of the arguments for and against each proposed solution idea

11.1.6.1. IS-VIOLATED-BY **missing arguments**
An idea is missing some important arguments for or against it.

11.1.6.1.1. IS-CAUSED-BY **self-focused**
The participants generate arguments that refer to criteria that concern them personally, but not those that impact welfare for other groups.

11.1.6.1.2. IS-HANDLED-BY **neglected criteria**
Few/no arguments have been created that evaluate a given idea with respect to one of the solution criteria for that problem.

11.1.6.1.3. IS-HANDLED-BY **few/no arguments**

i.e. there are (important) ideas that have few or no arguments attached to them

11.1.6.1.4. IS-HANDLED-BY **unbalanced arguments**
There is a large imbalance in the number of pros and cons in the debate over an idea.

11.1.6.1.5. IS-HANDLED-BY **few people contributed arguments**

11.1.6.1.6. IS-HANDLED-BY **idea/argument rating disconnect**
We can use such techniques as Bayesian inference (Bolstad, 2010) to propagate a user's ratings for arguments up the argument map to predict how the user should have rated the ideas these arguments address. If there is a large divergence between a user's predicted and actual ratings for an idea, that suggests that the user has not yet entered arguments that are compelling to him or her. A suggestion is to show users the possible misalignment between popularity ratings assigned freely by users and computed scores based on the structure of arguments and how much underlying pros and cons are supported. The gap can be used also as an incentive to users: a popular idea with poor computed score should invite its supporters to provide new arguments or improve the existing ones.

11.1.7. HAS-PART **high-quality argumentation**
The arguments entered in the summary map are well-founded.

11.1.7.1. IS-VIOLATED-BY **false premises**
the arguments made are based on false premises

11.1.7.1.1. IS-CAUSED-BY **argument sabotage**
Someone who disagrees with an argument edits it to sabotage it, rather than (properly) simply downrating or arguing against it.

11.1.7.1.1.1. IS-HANDLED-BY **argument edit wars**
assess the prevalence of edit wars (rapid alternating rollbacks) in the post edit histories – esp. by people who take differing positions on the issues they are editing arguments for. see ZIF workshop paper by János Kertész, Budapest: Edit wars on the Wikipedia:. an interesting measure for wikipedia article controversiality = # mutual reverts by more established user accounts (discounting young vandals)

11.1.7.1.2. IS-HANDLED-BY **low argument ratings**
An argument got a low average rating from the community.

11.1.7.1.3. IS-HANDLED-BY **expert evaluation**
Expert(s) judged that the argument is ill-founded.

11.1.7.2. IS-VIOLATED-BY **incorrect inference**
the arguments made are based on logical fallacies

11.1.7.2.1. IS-HANDLED-BY **automated feedback**
The use of artificial intelligence techniques holds promise to increase the effectiveness of argumentation systems by automatically analyzing user actions and providing supportive feedback. see: http//www.ascilite.org.au/ajet/ajet25/butchart.pdf Mclaren, B. M., Scheuer, O., & Mikšátko, J. (2010). Supporting collaborative learning and e-Discussions using artificial intelligence techniques. International Journal of Artificial Intelligence in Education, 20(1)(1), 1-46. Scheuer, O., McLaren, B. M., Loll, F., & Pinkwart, N. (2012). Automated Analysis and Feedback Techniques to Support Argumentation: A Survey. In: N. Pinkwart, & B. M. McLaren (Eds.), Educational Technologies for Teaching Argumentation Skills (pp. 71–124). Bentham Science Publishers.

11.1.7.2.2. IS-HANDLED-BY **expert judgment**
Expert(s) judge that the inference is faulty.

11.1.8. HAS-PART **high-quality ratings**
The community ratings for the argument map contents (i.e. whether issues are relevant, criteria are important, ideas are promising, arguments are compelling) are accurate.

11.1.8.1. IS-VIOLATED-BY **too few ratings**
There are too few ratings for a post to draw reliable conclusions about how the community judges it's value.

11.1.8.2. IS-VIOLATED-BY **dishonest ratings**
Ratings are dishonest.

11.1.8.2.1. IS-HANDLED-BY **rating inconsistency**
A participant gives inconsistent ratings e.g. they rate the arguments supporting an idea highly, but give the idea itself a poor rating.

11.1.8.3. IS-VIOLATED-BY **incorrect ratings**
the user's ratings for the posts are incorrect

11.1.8.3.1. IS-HANDLED-BY **missing suppoprt**
someone gave a strong + or - rating without a backup argument that they authored or highly rated. This can take several levels, according to the Discourse Quality Index (Jurgen Steiner): (1) The speaker does not present any arguments (asks, for example, merely for additional information) (2) The speaker only says that X should or should not be done, that it is a wonderful or a terrible idea, etc.. But no reason is given for why X should or should not be done. (3) The speaker justifies only with illustrations why X should or should not be done. (4) The speaker gives a reason Y why X should or should not be done. But no linkage is made why Y will contribute to X. (5) The speaker gives a reason Y why X should or should not be done, and a linkage is made why Y will contribute to X. (6) The speaker gives at least two reasons why X should be done and for at least reasons a linkage is made with X.

11.1.8.3.2. IS-HANDLED-BY **ignored arguments**
User did not attend to (read, rate) arguments when rating posts impacted by these arguments

11.1.8.3.3. IS-HANDLED-BY **irrational ratings**
see how well a model of rational rating predicts user's ratings. We can use knowledge about map structure to infer what a user's rating for a post "should" be given their ratings for posts below it in tree. For example, with mutex ideas, a rate for X is probably a rate against alternative to X.

11.2. REQUIRES **complete evaluation**
All the (promising) ideas are evaluated.

11.2.1. IS-VIOLATED-BY **evaluation groupthink**
everybody quickly converges to evaluating a very small set of ideas for an issue, ignoring the rest

11.2.1.1. IS-HANDLED-BY **attention narrowing**
we can measure when one idea under an issue is receiving the bulk of the community's argumentation and rating while competing ideas are neglected

12. HAS-PART **P: select best solution**

12.1. REQUIRES **outcome is broadly accepted**

12.1.1. IS-VIOLATED-BY **divisive issues**

12.1.1.1. IS-HANDLED-BY **highlight cross-cutting arguments**
Highlight existing arguments that appeal across balkanized groups in order to help develop increased consensus.

12.1.1.2. IS-HANDLED-BY **high pro/con activity**
do a histogram of activity level for different post types and see if pros and cons are unusually frequent

12.1.1.3. IS-HANDLED-BY **balkanizing issues**
we can use latent semantic indexing or principal components analysis or the like to find which are the issue sets that most divide people into clusters. A principal components analysis could help find, in effect, the fault lines in a debate, the sets of issues that most tend to divide people. We could find, for example, that abortion and gun control and school vouchers are highly divisive issues but if people agree on one of those issues they tend to agree on all the others well. We can then ask: what do these issues have in common? What underlying motivation or belief do they reflect? How can we attempt to reduce polarization along this dimension?

12.1.1.4. IS-HANDLED-BY **issues with high/growing rating variance**
where there are many arguments and contributors but no clear preponderance of highly-rated pros or cons

12.1.2. IS-VIOLATED-BY **many disaffected participants**
There are many deliberation participants who feel the selected outcome is unacceptable.

12.1.2.1. IS-HANDLED-BY **P: resolve: identify commonalities among participants**
cf Terry Steichen's work (the TopicCentral system) on finding commonalities in different people's favored portions of the deliberation map.

12.1.2.2. IS-HANDLED-BY **P: resolve: engage conflict resolution experts**
identified perhaps using analytics applied to deliberation summary?

12.2. REQUIRES **solution map is mature**
i.e. there is sufficient coverage of the issues, ideas, and arguments to make a decision

12.2.1. IS-ACHIEVED-BY **author/moderator activity dropoff**
If the fraction of author vs moderator contributions to a discussion drops, this suggests that the discussion is losing steam - it is only kept active by the effort of the moderators.

12.2.2. IS-ACHIEVED-BY **six hats**
Assess whether the problem solving session has progressed through a complete \"six hats\" program: Blue, White, Green, Red, Yellow, Black, which can be mapped to an argument map setting as follows:

- Thinking (Blue) - thinking about thinking, process issues
- Information: (White) - considering what information is available, what are the facts? provide details on the issues (= problems to be solved)
- Creativity (Green) - statements of provocation and investigation, seeing where a thought goes (= propose ideas)
- Good points judgment (Yellow) - logic applied to identifying benefits, seeking harmony (= identifying pros)
- Bad points judgment (Black) - logic applied to identifying flaws or barriers, seeking mismatch (= identifying cons)
- Emotions (Red) - instinctive gut reaction or statements of emotional feeling (but not any justification) (=> ratings? comments?) indicative of final stage - detected using prevalence of emotive words?

wikipedia article

12.2.3. IS-ACHIEVED-BY **"full" map topology**
map is both sufficiently bushy and deep.

12.2.4. IS-ACHIEVED-BY **completes narrative template**
The customer specifies the kind of narrative they want i.e. the main questions, the depth of argumentation, the breadth of options etc The system evaluates how far the argument map has gone to enabling that narrative, and asks the crowd to focus on the areas that yet need to be filled in. See work on rhetorical structures e.g. http//www.cs.columbia.edu/~kathy/NLP/ClassSlides/Slides09/Class20-Discourse/my-discourse.pdf

12.2.5. IS-ACHIEVED-BY **lifecycle stages**
There are many different possible models of the life stages a deliberation goes through as it matures. These include: • evolve from defining issues to proposing ideas to identifying increasingly broad and deep trees of pro and con arguments • evolve from creating new posts, to refining them, followed eventually by relative quiescence • opinion churn (i.e. whether the highest-rated ideas for individuals, as well as the community as a whole, are still changing rapidly or not) moderates as we reach the end of the lifecycle. • community support (as assessed by idea and arg ratings) concentrates on a few strongly supported ideas (lots of high ratings) • deliberation goes through the stages of preach to crowd, angry debunkers, filling in implicit support with reasoned data-based responses, irrelevant bored commentaryF • map growth tends to follow an S-shaped curve: map may be reaching maturity when slope decreases. http//crowdresearch.org/blog/?p=4602&utm_source=feedburner&utm_medium=email&utm_campaign=Feed%3A+FollowTheCrowd+%28Follow+the+Crowd%29

12.3. REQUIRES **high-quality votes**

i.e. they reflect the user's best judgment about which selection to make

12.3.1. HAS-PART **votes are truthful**

12.3.2. HAS-PART **votes are rational**
i.e. they represent a logically consistent response to all the relevant ideas and arguments (as opposed to some kind of bias)

12.3.2.1. IS-VIOLATED-BY **ignore higher-level context**
Users make decisions for low-level issues with taking into account higher-level issues and decisions that should have a major impact.

12.3.2.1.1. IS-HANDLED-BY **avoid: encourage hierarchical rating**
vote/rate on abstraction before voting on details

12.3.2.2. IS-VIOLATED-BY **hedgehog voter**
The voter has ignored ideas or arguments that should be relevant to the issue they voted on. We can call this the "hedgehog" exception, after Philip Tetlock, who pointed out that some people ("hedgehogs") only pay attention to a subset of information that is close to their original point of view (see Zaller), while other roam the information space more broadly ("foxes").

12.3.2.2.1. IS-HANDLED-BY **opinion shift**
If a user has substantially changed his ratings about ideas and arguments, this suggests he/she is open to being influenced by new information and perspectives.

12.3.2.2.2. IS-HANDLED-BY **user saw relevant posts**
Check if user has not read or rated relevant ideas and arguments under the issue he/she is voting on.

12.3.2.3. IS-VIOLATED-BY **voting cascades**
It has been shown that when people are asked to rate competing ideas, if they can see the ratings made to date (e.g. they see the ideas in popularity-sorted order), then the first ideas that happen to get a rating advantage tend to become the eventual winners—they “lock in” to the winning position—even if they are worse than ideas that appeared later or started with lower ratings (Salganik et al., 2006). its is therefore a problem if people vote for ideas based on their popularity (i.e. based on how many other people have voted for them) rather than their inherent merits.

12.3.2.3.1. IS-HANDLED-BY **ratings lock**
check whether the popularity order for a set of competing ideas remains relatively unchanged as the deliberation progresses

12.3.2.4. IS-VIOLATED-BY **bias**
participants is biased towards a given decision irregardless of arguments and other alternatives

12.3.2.4.1. IS-HANDLED-BY **motivated position change**
We can assess degree of bas by measuring whether users change their position through the course of the deliberation, or not. This can have several levels, according to Jurgen Steiner's Discourse Quality Index: (1) The speaker indicates a change of position. Gives as reason for change arguments heard during the experiment. (2) The speaker indicates a change of position. Does not refer to arguments heard during the experiment. (3) The speaker does not indicate a change in position. But does acknowledge the value of other positions heard during the experiment. (4) The speaker does not indicate a change of position. And does not acknowledge the value of other positions heard during the experiment.

12.3.2.4.2. IS-HANDLED-BY **coherence theory**
Can we use coherence theory, applied to the posts that the user ranked highly and thus presumably used in their decision, to assess the logical coherence of their votes? See http://www.iiia.csic.es/~joseph/index2.html

12.3.2.4.3. IS-HANDLED-BY **rating disconnect**
Assuming all the key arguments have be entered i.e. the map is mature: the user selects an idea that doesn't make sense given the ratings he/she gave to their underlying arguments and the competing ideas.

12.3.3. HAS-PART **votes are truthful**

12.3.4. HAS-PART **votes are rational**
i.e. they represent a logically consistent response to all the relevant ideas and arguments (as opposed to some kind of bias)

12.3.4.1. IS-VIOLATED-BY **ignore higher-level context**
Users make decisions for low-level issues with taking into account higher-level issues and decisions that should have a major impact.

12.3.4.1.1. IS-HANDLED-BY **avoid: encourage hierarchical rating**
vote/rate on abstraction before voting on details

12.3.4.2. IS-VIOLATED-BY **hedgehog voter**
The voter has ignored ideas or arguments that should be relevant to the issue they voted on. We can call this the "hedgehog" exception, after Philip Tetlock, who pointed out that some people ("hedgehogs") only pay attention to a subset of information that is close to their original point of view (see Zaller), while other roam the information space more broadly ("foxes").

12.3.4.2.1. IS-HANDLED-BY **opinion shift**
If a user has substantially changed his ratings about ideas and arguments, this suggests he/she is open to being influenced by new information and perspectives.

12.3.4.2.2. IS-HANDLED-BY **user saw relevant posts**
Check if user has not read or rated relevant ideas and arguments under the issue he/she is voting on.

12.3.4.3. IS-VIOLATED-BY **voting cascades**
It has been shown that when people are asked to rate competing ideas, if they can see the ratings made to date (e.g. they see the ideas in popularity-sorted order), then the first ideas that happen to get a rating advantage tend to become the eventual winners—they “lock in” to the winning position—even if they are worse than ideas that appeared later or started with lower ratings (Salganik et al., 2006). its is therefore a problem if people vote for ideas based on their popularity (i.e. based on how many other people have voted for them) rather than their inherent merits.

12.3.4.3.1. IS-HANDLED-BY **ratings lock**
check whether the popularity order for a set of competing ideas remains relatively unchanged as the deliberation progresses

12.3.4.4. IS-VIOLATED-BY **bias**
participants is biased towards a given decision irregardless of arguments and other alternatives

12.3.4.4.1. IS-HANDLED-BY **motivated position change**
We can assess degree of bas by measuring whether users change their position through the course of the deliberation, or not. This can have several levels, according to Jurgen Steiner's Discourse Quality Index: (1) The speaker indicates a change of position. Gives as reason for change arguments heard during the experiment. (2) The speaker indicates a change of position. Does not refer to arguments heard during the experiment. (3) The speaker does not indicate a change in position. But does acknowledge the value of other positions heard during the experiment. (4) The speaker does not indicate a change of position. And does not acknowledge the value of other positions heard during the experiment.

12.3.4.4.2. IS-HANDLED-BY **coherence theory**
Can we use coherence theory, applied to the posts that the user ranked highly and thus presumably used in their decision, to assess the logical coherence of their votes? See http://www.iiia.csic.es/~joseph/index2.html

12.3.4.4.3. IS-HANDLED-BY **rating disconnect**
Assuming all the key arguments have be entered i.e. the map is mature: the user selects an idea that doesn't make sense given the ratings he/she gave to their underlying arguments and the competing ideas.

12.4. REQUIRES **sufficient votes**
sufficient votes are available to fully, and fairly, capture the wisdom and preferences of the voters.

12.4.1. IS-VIOLATED-BY **insufficient votes**
There is insufficient preference information (in terms of votes or ratings) to pick a clear winner among the solution ideas.

12.4.1.1. IS-HANDLED-BY **confidence analysis**
Can we do an analysis to determine which ideas need to be assessed more completely in order to allow high-confidence selections of top-level solution ideas? Would we need some kind of confidence interval analysis? Could this be calculated based just on simple ratings, or would people need to express confidence scores for their ratings (e.g. very sure, not very sure).It would make sense to take into account the controversiality of the ideas e.g. if an idea is controversial, we would probably want to get more ratings for it to be more sure that people really prefer it (or not).

12.5. REQUIRES **votes aggregated properly**

12.5.1. HAS-PART **representative**
i.e. the decision created by aggregating the group's vote should represent what most individuals wanted

12.5.2. HAS-PART **fair**

12.5.3. HAS-PART **representative**
i.e. the decision created by aggregating the group's vote should represent what most individuals wanted

12.5.4. HAS-PART **fair**

13. HAS-PART **Enact solution**

13.1. REQUIRES **feasible solution**
The solution is a feasible one (i.e. can be implemented).

13.2. REQUIRES **commitment to action**
In social media, participation is high but incitement to action is historically low. Online debate and deliberation tools are populate by enthusiasts who have interest in the subject, spend time and efforts into debating it, but have not yet committed into taking action. How do we engage enthusiast/motivated audiences to translate the emerging trends and patterns into concrete actions to lead to further change?