CHAPTER XX

Measuring and Identifying Culture

Michael D. Young

Social Science Automation

Culture is widely regarded within academia and within the Intelligence Community (IC) as one of the most important factors that influence and predict human behavior. Agencies such as the Defense Intelligence Agency (DIA) are searching for ways to enable their analysts to identify, compare, and contrast the cultures of various populations in order to support planners and warfighters. For example, DIA Information Operations planners would like to know what the beliefs are that lead to support for the Taliban in Afghanistan, who holds those beliefs, and how those beliefs can be changed.

Most of the current approaches to measuring culture in the IC (HSCB 2009) are compatible with a propositional approach to culture that defines culture as:

> shared cognition (beliefs, values, attributes) that is transmittable between individuals and therefore across time and space.

The assumptions of this approach are that members of a culture will share and express the same beliefs, desires, and intentions and may also use their own vocabulary of concepts when they express those beliefs. Therefore, a culture can be identified as a set of individuals whose expressed beliefs and ways of expressing those beliefs are similar to one another's. For example, although we know that abortion is salient in both pro-life documents and in pro-choice documents, we expect the propositions about abortion to be quite different in each set of documents. All propositional approaches to culture require (A) the extraction of propositions from culture repositories such as speeches, newspapers, books, expert judgments, field reports, internet media, and other sources, and (B) subsequent analysis using network analysis techniques variously referred to as causal loop, influence, or cultural network analysis.

In previous work, Social Science Automation has used the Text Mapping coding scheme, Profiler Plus, and WorldView to apply the propositional approach to a set of advocacy documents (N = 78) written by opponents and proponents of gay marriage and abortion rights. The primary finding of those efforts was that overall document difference measures (transformation cost and incongruence) are insufficient to group documents and isolate cognitions that reflect a culture. The overwhelming "noise" that occurs naturally in unstructured texts produces difference scores in excess of 0.95 on a scale of 0.0 to 1.0. However, despite the low signal to noise ratio in unstructured texts from the "wild", several possibilities were identified to greatly reduce the noise. We have subsequently made progress on noise reduction, added DIA's Critical Network Analysis Tool (CNAT) to our toolkit, and identified several promising indicators that may, in combination, identify documents from a single culture.

1. PROMISING INDICATORS

Although the culture signal in our test documents is evident to human readers, the low signal to noise ratio in unstructured text appears to preclude any prospect of using transformation cost or other gross measures of difference to discriminate between cultures. However, the promising indicators explored in the current work are all intended to isolate or amplify aspects of the signal including:

- Shared concepts.
- Shared propositions.
- Concept valence.
- Threat/Victim.

1.1 Shared Concepts

A standard method of document clustering is to group documents based on the terms they share and on the co-occurrence of those terms (n-gram analysis, see also latent semantic analysis). However, this procedure does not retain the propositional content of the documents. Shared concept analysis is in some ways less powerful than n-gram analysis, but it can be performed within WorldView while retaining the propositional content of the documents.

WorldView does not provide any way to evaluate whether shared concepts analysis discriminates between groups. However, the shared concept and relation reports from WorldView can be transferred to CNAT which has routines for identifying cohesive groups. In CNAT, documents from the same culture (abortion versus same-sex marriage) are expected to form a strong cohesive subgroup that is distinguishable from the opposing culture with a classification accuracy significantly greater than 0.50 (chance). Although a plot in CNAT of the network of the 15 most shared concepts and their containing documents produces no useful discrimination between the two sets of documents, the discriminating power of each shared concept can be evaluated by examining changes in the non-directional cohesive strength for two groups as concepts are added and removed from the network and it should be possible to determine a minimum set of shared concepts which maximizes cohesive strength. If shared concepts do discriminate between the two cultures, the classification accuracy of groupings selected by CNAT should be greater than 0.5 (chance) on scale of 0.0 to 1.0

As a test of this approach, the cohesive strength measure was used to select the six concepts (marriage, rights, same-sex, abortion, couples, gay) that provided the best discrimination between the groups. When CNAT is allowed to select the two groups, an overall cohesive strength of 0.821 is achieved with a classification accuracy of 0.81. However, if documents sitting between the two groups are classified by hand, the cohesive strength increases slightly to 0.823 and classification accuracy increases to 0.91 (Figure 1.). This suggests that it is possible to construct an optimization algorithm that is blind to the meaning of the concepts and yet can still achieve high classification accuracy.

Figure 1. Connections between the six best most shared concepts and the same-sex marriage and abortion documents (N=78).

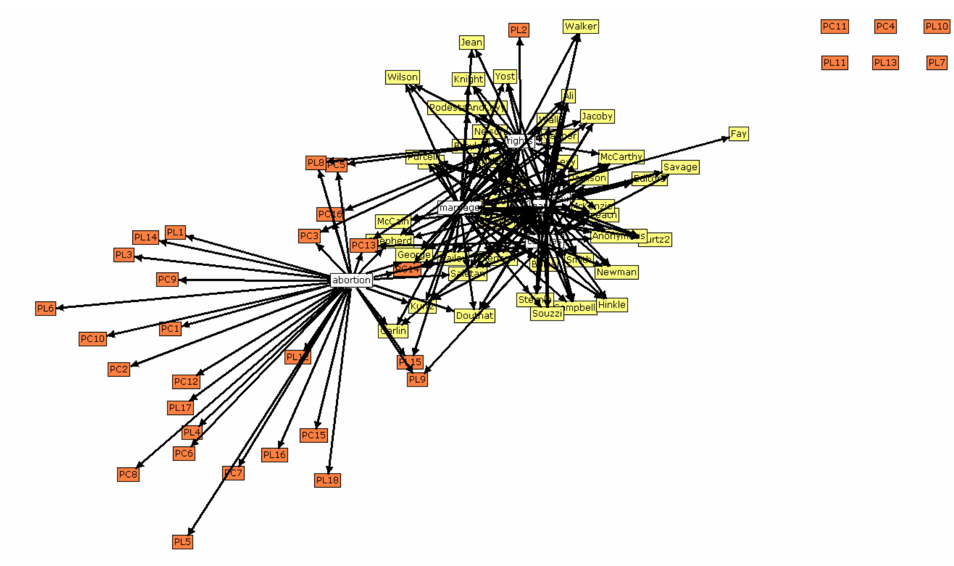

As a further test, a similar optimization methodology was repeated for each of the two document subsets (abortion, same-sex marriage). For the abortion document subset, two groups are obtained with 0.75 non-directional cohesive strength and a classification accuracy of 0.64 (Figure 2.). However, for the same-sex marriage document subset, only one group is obtained and the classification accuracy is barely above chance at 0.52.

Figure 2. Optimized groups for the abortion document set. Non-directional cohesive strength = 0.75, accuracy = 0.64

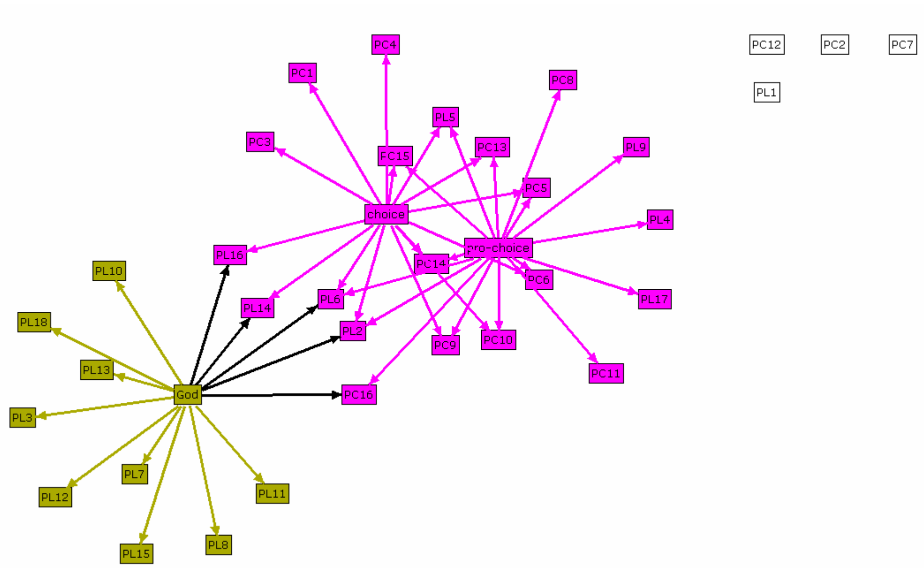

1.2 Shared Propositions

Although we expect cultures to differ in their use of concepts, some distinct cultures may use the same concepts, but in different ways. For example, abortion is salient in both the pro-life documents and in the pro-choice documents. However, the propositions about abortion should be quite different. Pro-life documents are likely to describe abortion as *wrong*. On the other hand, Pro-choice advocates may describe abortion as *a medical procedure*. Thus, although the concept abortion may be shared across the two cultures, the cognitions about that concept in each culture may be very different.

Repeating the analysis process used for shared concepts, a shared proposition analysis for the entire document set nicely separates the abortion documents from same-sex marriage documents, but unfortunately, it only yields one group and numerous singletons and both the cohesive strength and classification accuracy are undefined in this case. Applying the analysis process to the same-sex marriage and abortion documents separately did not yield any result with a classification accuracy greater than chance.

1.3 Concept Valence

In the absence of shared concepts or propositions that distinguish documents and cultures, it may be possible to distinguish documents and cultures by how they evaluate shared concepts. For example, using the nonsense word deetchzeeb, consider the following propositions taken from one of six imaginary documents: deetchzeeb is wrong, deetchzeeb is good, deetchzeeb is evil, deetchzeeb is honorable, deetchzeeb is criminal, and deetchzeeb is pleasant. All six propositions and imaginary documents share the concept deetchzeeb but none of them share a proposition. However, it is quite easy to group the six propositions (and their documents) into two groups, one where deetchzeeb is positive and one where deetchzeeb is negative.

To explore the usefulness of concept valence, a very rough concept valence prototype coding scheme was created and used to generate data for both sets of documents. The concept valence coding scheme simply tags words with an evaluation. For example:

> Abortion is wrong.
> ->
> (abortion valence1 true factual present bad1)

The initial results of a shared proposition analysis using valence data are promising (Figure 5.), but the classification accuracy of 0.58 is only a little better than chance. Although substantial work remains to improve the performance of the concept valence coding scheme, the initial results justify additional work.

1.4 Threat/Victim

One additional aspect of concept valence that can be explored is the directionality of the valence. For example, if we use a simple non-directional good/ bad valence indicator, "A attacks B" is "bad" for both A and B and this becomes a shared valence relation, tending to group A and B together. However, if the "threat" and "victim" of the proposition are distinguished they may provide greater discrimination. In both the same-sex marriage and abortion documents there is clear disagreement between the sides about what the threats are and who the victims are. For example, there is a prevalent belief among those opposed to same-sex marriage that "same-sex marriage will undercut marriage" even as some proponents believe that gay and lesbian Americans are discriminated against under current laws. Distinctions between threats and victims such as these are also likely to be valid and relevant in cultures of interest to the IC and may provide insight into cultural perceptions of constraints and threats. Complementary Benefit/Friend indicators may also prove useful but were not explored. For example, in many of the Anti same-sex marriage documents, the people, as voters in referenda, are seen as trustworthy.

In the initial exploration of Threat/Victim, a Victim concept (typically an actor) is any concept that is under attack or otherwise threatened; a Threat concept (also typically an actor) is a concept that is wielding illegitimate power or actively threatening or displaying a threatening posture. For example, in the hand-coded sentences below, examples of concepts that are coded threat are **bold** and concepts coded victim are *italic*. The terms that indicate these relationships are <u>underlined</u>.

- An even more substantive <u>danger lies in the consequences</u> of **gay marriage** on *the next generation.*
- There are valid -- and secular -- reasons to believe that **same-sex marriage** <u>will undercut</u> *marriage* itself.
- *Gay men and lesbians* <u>suffer</u> **discrimination**.

A prototype Threat/Victim coding scheme was developed that identifies threat and victim noun phrases. Illustrative Profiler Plus output for the sample sentences above is given below:

An even more substantive danger lies in the consequences of gay marriage on the next generation.
->
(gay power true attribute na (marriage power true na na marriage))
(next victim true attribute na (generation victim true na na generation))

There are valid -- and secular -- reasons to believe that same-sex marriage will undercut marriage itself.
->
((both same sex) power true attribute na (marriage power true na na marriage))
(marriage victim true na na marriage)

Gay men and lesbians suffer discrimination.
->
(gay victim True attribute na (man victim true na na man))
(lesbian victim true na na lesbian)

A shared proposition analysis of the Threat/Victim data produces two groups with an overall non-directional cohesive strength of 0.822 and a classification accuracy of 0.68 (Figure 6.) providing grounds for further analysis.

A subsequent shared proposition analysis applied to the 34 same-sex marriage documents with clear positions produced two groups with non-directional cohesive strength of 0.888 and a classification accuracy of 0.71. These results provide further evidence both that creating a culture identification routine is possible, and that combining data from more than one indicator may produce increased classification accuracy. However, applying the same procedure to the abortion documents proved less successful with a classification accuracy less than chance (0.41).

1.5 Combined Indicators

Three of the indicators examined provide some evidence of discriminatory power and, although there is substantial overlap between document groups identified for each indicator the overlap is not complete. This partial overlap suggests that the indicators may perform even better in combination. To assess this possibility, a shared item analysis was conducted for all 78 documents using both the 9 best concepts and 7 best valenced concepts producing two groups with a non-directional cohesive strength of 0.73 and a classification accuracy of 0.95 (Figure 3.) Unfortunately this success was not repeated with the addition of the Threat/Victim data.

Figure 3. Best 9 Concepts and 7 valenced concepts for all documents. Non-directional cohesive strength = 0.73; classification accuracy = 0.95.

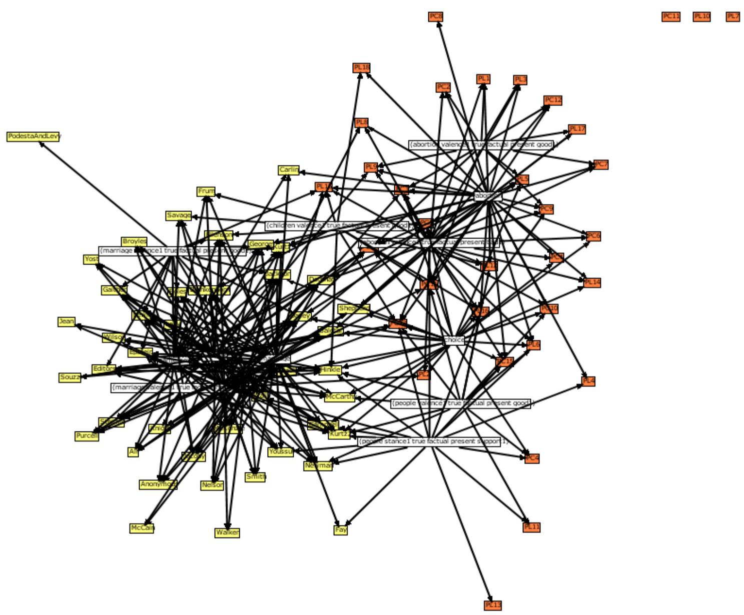

1.6 Conclusion

Despite a discouraging start to our investigations, our work to date suggests that, with some refinement, the propositional approach will lead to a useful methodology for identifying cultural groups and measuring cultural content.

REFERENCES

HSCB OSD Human Social Culture Behavior Modeling Program Focus 2010 Conference, 5-7 August 2009. See in particular:

- Modeling Populations in Stability Operations Modeling Populations in Stability Operations using Cultural Geography, U.S. Army TRADOC Analysis Center—Monterey
- UPenn’s Leader Modeling Capability:An Holistic Socio-Cognitive Agent Approach, Dr. Barry G. Silverman
- An Architecture for Interpersonal Interactions; Politeness Modeling and Where it Can Take You, Christopher A. Miller & Peggy Wu.
- Web-Based Geospatial Visualization Web Based Geospatial Visualization for HSCB Behaviors, William Wright, David Schroh, & David Jonker.
- Plug and Play Avatars for Training and Mission Rehearsal, Ed Sims & Keith Knudsen.
- The First-Person Cultural Trainer, The University of Texas at Dallas.
- Sociocultural Communication Skills Training – A Serious Game Approach, Ali Ahmad & Colleen Phillips.
- A Cognigraphic Training Tool For Human Terrain Understanding, Celestine A. Ntuen.
- PSTK: A Toolkit for Modeling and Simulation of Power Structure Networks, Glenn Taylor
- The Environment ‐ Attitudes ‐ Actions Theory of Human Behavior, Dr. Van Parunak
- Modeling Strategic Contexts: The Axes of Uncertainty, David L. Sallach
- Lessons Learned from the Field, Martin Kruger
- Cultural Network Analysis: Method and Application, Winston R. Sieck
- *Conceptualizing our Nations Together: A Cultural Taxonomy (CONTACT)*, Denise Nicholson, Eduardo Salas, Randy Shumaker, Shawn Burke, Daniel Barbar, & Maritza Salazar.
- *Training and Mission Rehearsal*, Allison Abbe.
- *Temperament as a Lens for Culture*, Sae Schatz, Denise Nicholson & Marissa L Shuffler.
- *Cross-cultural Analysis Methods Using Advertising and Psychophysiological Measures*, Cali Fidopiastis, Stephanie Lackey, Denise Nicholson.