

The Methodological Challenges of Extracting Dark Networks: Minimizing False Positives through Ethnography

Michael Kenney & Stephen Coulthart

Incomplete information pervades the study of dark networks. Illicit entrepreneurs' desire for secrecy, to protect the confidentiality of their affairs and associations, means that many links and nodes remain hidden from observation. Those determined to observe such networks, whether law enforcers, journalists, or researchers, generally do so as unwanted interlopers, compelling them to make uncertain inferences based on imperfect knowledge. The Department of Defense and intelligence community have tried to address this problem by investing millions of dollars in computational tools that trawl through enormous amounts of documents, identifying and coding covert networks from secondary sources, such as news reports and court records. Yet the validity of the networks extracted with automated tools remains uncertain given problems with the accuracy of the data that goes into the models. Scholars using these tools confront similar challenges in compiling reliable network data from incomplete information.

This chapter describes our efforts to understand one dark network, al-Muhajiroun, an outlawed activist network in Britain, through computational network analysis of a dynamic, yet flawed source of data: news reports. After compiling thousands of news reports on al-Muhajiroun into a single data set, we used *AutoMap* and *Organizational Risk Analyzer* (ORA), two software tools utilized by intelligence analysts, to extract and examine social networks from these data. As helpful as the tools were, we experienced a number of challenges in using them, challenges we document in this chapter. Our focus is methodological, rather than empirical

This research was supported by a grant from the Office of Naval Research, U.S. Department of the Navy, No. N00014-09-1-0667. Any opinions, findings, or recommendations expressed in this chapter are those of the authors and do not reflect the views of the Office of Naval Research. The authors wish to thank John Horgan, Kathleen Carley, Michael Bigrigg, and other members of the CASOS team for their help during the research project. We alone are responsible for any errors.

or analytical. We discuss our experience using AutoMap and ORA to identify lessons learned for other practitioners of dark network analysis, rather than test formal hypotheses or develop causal explanations of network behavior.

We offer these observations not as computer scientists who understand the algorithms and routines that run the software, but as ethnographers and social scientists steeped in Islamist militancy. Since 2008, we have engaged in extensive fieldwork on al-Muhajiroun, interviewing and observing dozens of participants at different events, including political protests, *da'wah* stalls, and study circles. Our immersion into the world of al-Muhajiroun, acquired gradually over many months, gave us inside knowledge on this collection of activists, knowledge we drew on to analyze the social networks extracted from the newspaper data. While traditional ethnography is often devalued in intelligence circles because of its slow pace and lack of scalability, our experience suggests that it offers an essential complement to computational network analysis of secondary-source data.

I. Researching al-Muhajiroun

Al-Muhajiroun is a transnational advocacy network that seeks to establish the global caliphate within and outside Britain through a combination of preaching (*da'wah*) and political activism. Since 1996, al-Muhajiroun activists have sought to spread their fundamentalist interpretation of Islamic scripture by organizing *da'wah* stalls on busy thoroughfares in numerous towns and cities, convening public conferences addressing a range of political and religious topics, and engaging in high-profile demonstrations designed to attract widespread media coverage. While al-Muhajiroun specializes in high-risk activism, pushing the bounds of legally protected speech and association through its provocative meetings and protests, over the years network activists have been implicated in numerous acts of political violence, including a suicide bombing that killed three people and injured fifty-five more at a popular bar in Tel Aviv in 2003, and, more recently, the brutal slaying of an off-duty British soldier on the streets of London in May 2013 (Wiktorowicz, 2005; Simcox, Stuart, & Ahmed, 2010; Burns & Cowell, 2013).

Beginning in the late 1990s, and accelerating after 9/11, al-Muhajiroun came under increasing pressure from government authorities. With the introduction of new counterterrorism legislation, previously legal acts from al-Muhajiroun's repertoire of contentious action, including raising funds for overseas fighters and "glorifying" terrorist violence, became illegal. The government also pressured the group's founder and spiritual leader, Omar Bakri Mohammed, to leave Britain, convicted several

prominent activists for a variety of “terrorism-related” offenses, and made it a crime to belong to the activist network by banning al-Muhajiroun and several spin-off groups. Rather than destroying the network, authorities merely drove it underground, as activists continued their illegal efforts under the covert guise of new spin-off groups (Raymond, 2010; Kenney et al., 2013).

Over a period of several years, we gathered extensive data on this dark network through firsthand fieldwork and secondary-source research. The purpose of this research was to combine ethnographic fieldwork and computational modeling to improve our understanding of how extremist networks learn from information and experience and how they adapt in response to external pressure. Between 2008 and 2013, we conducted five months of field research in London over the course of five separate trips. During these research trips we recorded more than 100 interviews with eighty-six respondents. Forty-seven of these respondents were current al-Muhajiroun activists, ranging from leaders to rank-and-file participants, many of whom we interviewed more than once. Apart from these interviews, we also spent considerable time observing activists at numerous demonstrations, *da‘wah* stalls, and public talks, and hanging out with them at restaurants and their study center in East London. We estimate that we spent between 100 and 200 hours engaged in face-to-face contact with al-Muhajiroun activists during our ethnographic fieldwork.

Over the same period, we also spent many additional hours collecting secondary-source data on al-Muhajiroun, primarily from English-language newspaper articles on the Islamist network in Lexis-Nexis. After screening for duplicates and unrelated articles, we ended up with a final data set of 3,306 newspaper articles published in sixty-four different newspapers. Dates of publication ranged from 1996, the year al-Muhajiroun was formed in Britain, through November 2012, the end point for our secondary-source data collection. We used AutoMap and ORA to extract and analyze our social network of al-Muhajiroun from these news reports.

II. Network Text Analysis and AutoMap

Since the 9/11 terrorist attacks, researchers and security officials have shown increasing interest in modeling dark networks (e.g., Xu & Chen, 2008; Morselli, 2009; Helfstein & Wright, 2011; Everton 2012b). In computational studies, dark networks are often modeled using network text analysis (Diesner & Carley, 2005; Graham, Carley, & Cukor, 2008; Carley, Bigrigg, & Diallo, 2012). This is a semi-automated method for mining networks from a variety of texts, ranging from news articles to

Table 4.1. *Example from al-Muhajiroun thesaurus*

Associated concepts	Root concepts	Attributes
Omar Bakri Mahammed	omar_bakri	Agent
Omar Bakri Mohammad	omar_bakri	Agent
Omar Bakri Mohammed	omar_bakri	Agent
Omar Bakri Mohammad	omar_bakri	Agent
Omar Bakri Mohammud	omar_bakri	Agent
Omar Bakri Muhammad	omar_bakri	Agent
Omar Bakri Muhammed	omar_bakri	Agent
Omar Bakri Muhammod	omar_bakri	Agent
Omar Bekri Mohammed	omar_bakri	Agent
Omar Bakri Fostock	omar_bakri	Agent
Omar Bakri Mahamed	omar_bakri	Agent
Omar Bakri Mohamad	omar_bakri	Agent
Omar Bakri Mohamed	omar_bakri	Agent
Omar Bakri Mohamud	omar_bakri	Agent
Omar Bakri Mohmmed	omar_bakri	Agent

intelligence reports. As Diesner points out, network text analysis extracts “meaning from texts by finding or establishing links between concepts and conducting network analysis of the resulting data” (2012, p. 83). Concepts are represented as nodes and connections are represented as ties between concepts that appear within a certain distance of each other in the text.

In this research we used AutoMap to perform network text analysis. Developed by the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University (Carley et al., 2011), AutoMap has become one of the most common semi-automated network text analysis programs used today. Intelligence analysts use the software to code and extract covert networks from large volumes of data, and researchers use it to investigate legally sanctioned “bright” networks (Kim, 2011; Lanham, Morgan, & Carley, 2011) and illicit “dark” networks (Carley, 2006a; Hutchins & Benham-Hutchins, 2011).

Using AutoMap to perform network text analysis requires the extraction of a “thesaurus,” or list of concepts, from the texts under analysis. The software uses this list to identify nodes in the data, similar to the role a codebook plays in hand coding. The thesaurus file, which analysts typically maintain as an Excel spreadsheet, contains three main columns: root concepts, associated concepts, and attributes. Root concepts form nodes in the network. As Table 4.1 demonstrates in reference to the global leader of al-Muhajiroun, roots can be paired with multiple associated concepts, which is necessary to reconcile conflicting spellings

or multiple aliases to a single node, thereby precluding the same person from appearing multiple times in the extracted network. Attributes, which appear in the third column of the example, identify the type of concept that is being extracted from the text. Because our research sought to understand changes in the al-Muhajiroun social network over time, we focused our thesaurus on agents, specifically individuals associated with al-Muhajiroun (e.g., Omar Bakri Mohammed). Consequently, the attributes column of our thesaurus contained only “agents.” However, when analysts seek to assess “meta-networks” containing multiple types of actors, the attribute column serves to distinguish between organizations, knowledge, resources, events, tasks, roles, locations, and agents (Carley, 2003).

III. Creating and Refining Thesauri

From the beginning of our research, our efforts to understand al-Muhajiroun were hampered by fuzzy boundaries and incomplete information. Many prominent individuals associated with the activist network are also connected to other Islamist groups and individuals who have received substantial attention in the media. Separating well-connected, committed activists from others who share al-Muhajiroun’s ideology without necessarily acting on its behalf was a major challenge, one intensified by biases in the newspaper data and activists’ desire to maintain the integrity of al-Muhajiroun’s operations in a hostile environment.

One of the main sources of data in our research was several thousand news reports on al-Muhajiroun, published from January 1996 through November 2012. While these reports contain valuable information about the activist network, many journalists mention prominent terrorists and Salafi-jihadi preachers in their accounts of al-Muhajiroun, irrespective of whether these individuals belong to al-Muhajiroun or support its activities operationally. For example, 340 news reports, representing just over 10 percent of the 3,306 articles in our data set, mention Osama bin Laden in the same article as Omar Bakri, al-Muhajiroun’s founder and spiritual leader.

After reading through a random sample of these articles, we determined that the links between Osama bin Laden and al-Muhajiroun were spurious. None of the articles focused on operational connections between Osama bin Laden and al-Muhajiroun. Instead, they discussed al-Muhajiroun’s attempts to build its brand by ideologically aligning itself with al-Qaida. Had any of the news reports we read discussed operational connections between al-Qaida and al-Muhajiroun, we would have “red flagged” Osama bin Laden and investigated the connections

through additional research, as we did for other individuals we flagged. In the few cases where our research suggested that these operational connections were accurate and reliable, we added these individuals to our list of al-Muhajiroun figures. In the case of bin Laden, however, the news reports suggested that the connections between him and al-Muhajiroun were ideological, not operational, so we did not flag him for additional analysis. Yet when these same articles mention bin Laden and other al-Qaida figures in their reporting on al-Muhajiroun, AutoMap identifies them as agents in the same network, creating false positives that distort the analysis.

News reports also typically fail to identify rank-and-file activists because of their desire to remain anonymous and protect the secrecy of their operations. When activists are mentioned in news reports, they are often identified by aliases or nicknames provided by the activists themselves, many of whom do not want their real names appearing in print. Activists also change their nicknames on occasion, further complicating outsiders' efforts to identify activists that do not wish to be identified. If reporters' tendency to highlight infamous figures like Osama bin Laden in their accounts of al-Muhajiroun increases the likelihood that false positives will undermine the results of networks extracted from these data, their failure to identify dozens of committed activists increases the chance that false negatives will do the same.

Our attempts to build thesauri from which we could extract accurate models of al-Muhajiroun's social networks were forced to confront these biases and omissions. While shortcomings in the news reports raise questions about the appropriateness of these data for our research, they were the only data available that we could use to capture social networks of al-Muhajiroun activists over time. Our interview and participant observation data were limited to the shorter period during which we conducted fieldwork, three years, as compared to the sixteen-year period the news reports covered. Given this constraint, we chose to work with the data we had, rather than abandon our efforts to map out al-Muhajiroun's networks over the longer period. One important advantage in using dynamic network analysis tools like AutoMap and ORA is that the networks extracted can be analyzed over time, allowing researchers to identify changes in network structure. But to do so, the networks must be partitioned into distinct time periods, which news reports allow, but interviews do not.

Instead of creating an al-Muhajiroun thesaurus from scratch, we sought to leverage CASOS' previous work by using its thesauri as the basic building blocks for our concept list. We hoped that using CASOS' thesauri would save us time, but we were surprised to learn that modifying the thesauri to our research project required substantial time and effort. We estimate that we spent approximately sixty man-hours

reworking the original CASOS thesauri into a preliminary thesaurus for our research on al-Muhajiroun.

We began by extracting the root concept names and associated concept aliases of forty-two unique al-Muhajiroun activists and affiliates from a larger “union” list of names maintained by CASOS. We added these entries into CASOS’ “universal” thesaurus to create a preliminary thesaurus for our research. This thesaurus was designed for meta-network analysis. Consequently it contains hundreds of thousands of concepts corresponding to locations, organizations, knowledge, resources, tasks, and events, as well as agents. We deleted many entries from this initial thesaurus, while adding a few new ones, gradually refining the generalized concept list to meet the needs of our research. Following several rounds of revisions, we produced a revised thesaurus containing 28,727 concepts, many of which still corresponded to concepts designed for meta-network analysis. Recognizing that many of these concepts were beyond the scope of our analysis, which sought to identify al-Muhajiroun’s agent-based social networks over time, we extracted an agent-based thesaurus from this larger revised thesaurus. This smaller, agent-based thesaurus contained 2,402 associated concepts corresponding to 364 root concepts, each representing a unique individual.¹

IV. Preliminary Analysis of Thesaurus

In our initial review of this thesaurus we encountered many agents who were affiliated with al-Qaida and other terrorist groups rather than al-Muhajiroun. This was not surprising given that one of the original sources of the thesaurus was CASOS’ Counter-Terrorism Database, which included entries for 585 individuals involved in terrorism (Gerdes & Carley, 2009). Having already invested substantial time and effort in modifying and extending this agent-based thesaurus, we decided to test its accuracy. We uploaded the agent-based thesaurus into the AutoMap software, using it to extract social networks from our newspaper data set. We partitioned these networks into three distinct periods, corresponding to important developments in al-Muhajiroun’s history.² We then

¹ Many individuals in the agent-based thesaurus had aliases and names with alternate spellings, accounting for the larger number of associated concepts.

² Network A contains the social network of al-Muhajiroun agents extracted from news reports published from the date of al-Muhajiroun’s founding in early 1996 through October 4, 2004, when the group voluntarily disbanded under government pressure. Network B contains agents in articles published beginning with the 7/7 London Tube and bus bombings in 2005, an event that precipitated Omar Bakri’s departure to Lebanon, to July 16, 2006, the day before the British government officially banned al-Muhajiroun’s successor organizations, al-Ghurabaa and the Saved Sect. Finally, Network C runs from July 17, 2006 to the end of the 2009 calendar year (Kenney et al., 2013).

performed social network analysis on these networks using ORA. The purpose of the analysis was to see how well our thesaurus worked with AutoMap and ORA in identifying key individuals from al-Muhajiroun.

We operationalized “key individuals” as the top ten ranked agents for two standard measures in social network analysis, betweenness centrality and eigenvector centrality. We chose these two measures because of their widespread application in social network analysis to identify brokers, people who connect different network nodes (betweenness centrality), and gatekeepers, people who link to highly connected nodes (eigenvector centrality) (Helfstein, 2012). Our attempt to identify top-ranked agents in al-Muhajiroun is similar to counterterrorism analysts’ network analyses, which seek to destroy terrorist networks by identifying and removing key individuals who hold the networks together (Roberts & Everton, 2011).

The results of our preliminary analyses were, at best, mixed. ORA did identify numerous al-Muhajiroun leaders among the top ten agents for betweenness and eigenvector centrality, including Omar Bakri, Anjem Choudary, Abu Izzadeen, and Abdul Rahman Saleem. Saleem’s high ranking in betweenness centrality was particularly encouraging. While he was not the most prominent al-Muhajiroun spokesman, Saleem was a long-standing broker in the network that connected activists in Britain and Pakistan. Unfortunately, our preliminary analysis failed to identify other prominent al-Muhajiroun activists among the top-ranked nodes, particularly in Network C, which corresponded to when we conducted our field research. More troubling, our analysis identified numerous false positives, among which we identified two types.

Nodal false positives refer to individuals in the top ten for betweenness or eigenvector centrality who, in our estimation, have never been activists or operational affiliates of al-Muhajiroun. These include prominent terrorists, such as Osama bin Laden and Richard Reid, and well-known Salafi-jihadi preachers, like Abu Hamza and Abu Qatada. These individuals are often mentioned in news reports on al-Muhajiroun, but lack operational connections to the activist network. Other nodal false positives include less well-known individuals implicated in various terrorist plots, but who are also not operationally linked to al-Muhajiroun (see Tables 4.2 and 4.3).

Relational false positives refer to individuals in the top ten for betweenness and eigenvector centrality who were operationally connected to al-Muhajiroun at some point, but who never attained the prominent positions in the network implied by their centrality rankings. These include Omar Sharif, one of the perpetrators behind the “Mike’s Bar” suicide bombing in Tel Aviv in 2003; along with Omar Khyam, Saladhuddin Amin, Waheed Mahmood, Jawad Akbar, and Mohammed Babar, all of whom were involved in a failed plot to bomb several targets in Britain

Table 4.2. *Betweenness centrality in al-Muhajiroun using preliminary thesaurus*

Rank	Network A	Between. Centrality	Network B	Between. Centrality	Network C	Between. Centrality
1	Abdul Saleem	0.042	Abdul Saleem	0.078	Abu Izzadeen	0.013
2	Osama bin Laden	0.037	Abu Hamza	0.060	Osama bin Laden	0.012
3	Abu Hamza	0.035	Abu Izzadeen	0.051	Anjem Choudary	0.011
4	Hassan Butt	0.020	Osama bin Laden	0.039	Abu Hamza	0.009
5	Saladhuddin Amin	0.013	Saladhuddin Amin	0.027	Mohammed Omran	0.009
6	Abu Qatada	0.011	Omar Bakri	0.020	Omar Khyam	0.008
7	Omar Sharif	0.011	Omar Sharif	0.018	Abdul Saleem	0.007
8	Waheed Mahmood	0.010	Abu Uzair	0.018	Abu Qatada	0.007
9	Mohammed Omran	0.010	Hassan Butt	0.016	Omar Bakri	0.005
10	Omar Bakri	0.009	Anjem Choudary	0.011	Mohammed Babar	0.005
	Nodal False Positive		Relational False Positive		AM Activists/Associates	

Table 4.3. *Eigenvector centrality in al-Muhajiroun using preliminary thesaurus*

Rank	Network A	Eigenvector Centrality	Network B	Eigenvector Centrality	Network C	Eigenvector Centrality
1	Abdul Kahar Kalam	1	Abdul Kahar Kalam	1	Anjem Choudary	1
2	Omar Sharif	1	Omar Bakri	1	Abdul Saleem	1
3	Abdul Karim	1	Richard Reid	1	Willie Brigitte	1
4	Younis al Hayyari	1	Abdul Karim	1	Omar Bakri	0.960
5	Abu Obeida	1	Younis al Hayyari	1	Saladhuddin Am in	0.918
6	Ramadan Shallah	1	Ezzit Raad	1	Jawad Akbar	0.787
7	Ezzit Raad	1	Fadal Sayad i	1	Omar Khyam	0.704
8	Abdul Koyair Kalam	1	Abdul Koyair Kalam	1	Waheed Mahmood	0.407
9	Abdul Qassim	1	Anjem Choudary	0.626	Abu Izzadeen	0.390
10	Mohammed Salim	1	Abu Izzadeen	0.562	Abu Hamza	0.343
	Nodal False Positive		Relational False Positive		AM Activists/Associates	

in 2004; and the Kalam brothers, Abdul Kahar and Abdul Koyair, whose house in East London was raided by British counterterrorism police in 2006 in a futile attempt to locate a suspected bomb factory (see Tables 4.2 and 4.3). All of these individuals were active in al-Muhajiroun, at least for brief periods, attending meetings and protests. In some cases, they met each other through these events. However, none of these individuals became prominent activists, let alone leaders in al-Muhajiroun, contrary to what their rankings suggest. Their high centrality rankings reflect the fact that the terrorist attacks or counterterrorism incidents with which they were associated were widely reported by the press and many of these reports mention al-Muhajiroun.

Nodal and relational false positives dominated the initial results of the al-Muhajiroun thesaurus. Of the twenty-eight unique individuals listed among the top ten agents for betweenness and eigenvector centrality in our preliminary analysis, thirteen were nodal false positives, representing 46 percent of the ranked agents. Another eight individuals were relational false positives, representing 29 percent of the top ten agents. In other words, 75 percent of the top-ranked agents for betweenness and eigenvector centrality were false positives, either nodal or relational. Conversely, only seven individuals out of twenty-seven, representing 25 percent of the top ten ranked agents, were key figures in al-Muhajiroun.

In other words, our preliminary thesaurus identified three times as many false positives as hits, not a satisfying result for an analysis meant to identify the brokers and gatekeepers of al-Muhajiroun. Any efforts to dismantle al-Muhajiroun based on these results would have done little damage to the network being targeted. The agent-based thesaurus clearly needed more work if we hoped to produce more accurate models of the activist network.

V. Refining the Thesaurus

Mindful of the false positives in our initial analysis, we set out to create the cleanest list of al-Muhajiroun activists and affiliates we could, hoping this would improve the accuracy of the networks AutoMap extracted from our newspaper data set. This required additional time and effort; over a period of four weeks, we spent approximately eighty man-hours refining the thesaurus. Thus, we invested a total of 140 man-hours in this document, to say nothing of the time the CASOS researchers invested in building the original list that served as our starting point.

In refining the thesaurus, we worked from our agent-based concept list, carefully reading through the spreadsheet line by line to determine which agents should stay and which should go. Drawing on our ethnographic

understanding of al-Muhajiroun, we conceptualized network activism and affiliation in operational terms. It was not enough for an individual to share an affinity with al-Muhajiroun's ideology or be involved in a widely reported terrorist plot to be defined as a node in the network. Rather, we sought to include only individuals who, to the best of our knowledge, had been operationally active in the network at some point in time. These were people who attended al-Muhajiroun's public lectures and private *halaqahs* (study groups) and who participated in the group's public protests and semi-private conferences. In applying these criteria we "red flagged" a host of al-Qaida leaders, Hizballah figures, participants in independent terrorist plots, and prominent Salafi-Jihadi preachers for additional scrutiny. Before removing any of these people from the thesaurus we conducted additional research on them to make sure that AutoMap had not identified network activists we did not know about. In most but not all cases, the red flags turned out to be false positives. The individuals in question were mentioned in news articles on al-Muhajiroun, but had few, if any, operational ties to the network. At the risk of engaging in confirmation bias, we chose to remove them from the thesaurus.

Our thesaurus cleaning was not limited to deleting entries. To address false negatives and the problem of missing agents, we also added several dozen individuals to the thesaurus: people we knew from our fieldwork and secondary-source research were active in al-Muhajiroun. Many of these names came from a separate coding file for qualitative analysis we built over several years that included individuals' real names, along with their known aliases. Integrating this file with the agent thesaurus produced dozens of new entries. We also discovered several duplicate entries, typically alternate spellings for the same individual that appeared as two different root concepts. We fixed these errors by merging the alternate spellings into a single root concept. Finally, we added alternate spellings for numerous al-Muhajiroun activists whose names reporters often misspell. After making all these changes to the thesaurus, we ended up with a final agent-based thesaurus of 2,047 entries for 337 unique individuals, an average of six different spellings and aliases for each agent.

VI. Improved Thesaurus, Better Results

After all our efforts to improve the thesaurus, were we able to create more accurate models of al-Muhajiroun's social networks? To find out we plugged the new and improved thesaurus into AutoMap to extract agent-based networks from the newspaper data. As in our first analysis,

we partitioned the networks into three distinct periods, corresponding to major developments in al-Muhajiroun's history.³ Like our initial analysis, we measured nodes in each network for betweenness centrality and eigenvector centrality. Once again, like counterterrorism intelligence analysts, we wanted to see how well our thesaurus worked in identifying key individuals from al-Muhajiroun.

The improvements in our final analysis were striking. Among the twenty-eight unique individuals ranked in the top ten for betweenness and eigenvector centrality, there was not a single nodal false positive. In all three networks, top-ranked agents for both measures were al-Muhajiroun activists or affiliates. By way of comparison, our preliminary analysis had included thirteen nodal false positives, which represented 46 percent of the top-ranked agents.

However, our final results still showed numerous *relational* false positives among the top-ranked agents. In fact, the number of relational false positives increased from eight in our preliminary analysis to nine in our final analysis, representing 32 percent of the agents for betweenness and eigenvector centrality in all three networks. As in our preliminary analysis, these included al-Muhajiroun activists and affiliates involved in terrorist plots that received a great deal of media coverage, inflating their centrality rankings in the models extracted from the newspaper data. For example, both Omar Sharif and Asif Hanif had reportedly been active in al-Muhajiroun prior to participating in the suicide bombing at "Mike's Bar" in Tel Aviv in 2003, but neither terrorist had been a central broker in the network, as their betweenness centrality scores suggested. Likewise, Omar Khyam and Mohammed Babar may have been key players in the foiled plot to bomb different targets around London in 2004, but they were not key players in al-Muhajiroun. Four relational false positives, including Afzar Munir, Aftab Manzoor, Mohamad Omar, and Yasir Khan, all went to Afghanistan in the heady days after 9/11,

³ However, we changed the partition dates in our analysis, reflecting our desire to measure the networks before and after al-Muhajiroun's spiritual leader, Omar Bakri, left Britain permanently, and before and after network leaders decided to relaunch the al-Muhajiroun platform in 2009. In the final analysis, Network A contains the social network of al-Muhajiroun agents extracted from news reports published from al-Muhajiroun's founding in early 1996 to the day before Omar Bakri left Britain for Lebanon, several weeks after the 7/7 bombings in London. Network B contains agents in articles published beginning the day Bakri left Britain to the day Anjem Choudary and other leaders officially relaunched the al-Muhajiroun platform in May 2009. Network C runs from the day after al-Muhajiroun's official rebirth to the end of our newspaper data collection in November 2012. While changing the partition dates altered individual agents' centrality scores and rankings, we do not believe it significantly impacted the number of false positives in the analysis, given that the false positives problem was largely due to the initial thesauri.

when al-Muhajiroun recruited fighters for jihad against British and American troops fighting there. While these individuals were reportedly involved in al-Muhajiroun, none of them were as central to the network as their rankings for betweenness and eigenvector centrality suggested. Finally, Ali Behesti participated in al-Muhajiroun protests in London, before leaving the group and firebombing the home of a British publisher who planned to release a controversial novel about Aisha, the Prophet Mohammed's wife and a venerated figure in Islamic culture.

Our revised thesaurus correctly identified all these individuals as al-Muhajiroun agents, but their centrality scores were inflated by their involvement in violent incidents that received extensive attention from the media. This reflects the poor quality of the news reports. Widespread, often sensational reports on terrorist attacks and other violent incidents linked to al-Muhajiroun increase the centrality scores of individuals mentioned in these accounts. Once we cleaned up the original concept list, the newspaper data itself became the problem.

On the positive side, the final results for betweenness and eigenvector centrality also identified numerous individuals we expected to see in a list of top-ranked agents, given their prominence in al-Muhajiroun. These included not only the al-Muhajiroun spokesmen who also appeared in the preliminary analysis, but several "false negative" activists we added when revising the thesaurus. We noticed the importance of these activists during our fieldwork, when we observed them interacting with other activists at different events. Some of these individuals were convicted of "terrorism-related" offenses in criminal trials that received substantial media coverage. Others were key figures in al-Muhajiroun spin-off groups that organized high-profile protests that received substantial media coverage. All of them were key players in the network, committed activists who essentially ran al-Muhajiroun.

In sum, our final analysis produced networks that more closely matched our expert conception of al-Muhajiroun, validating our efforts to refine the thesaurus. In our preliminary analysis, only 25 percent of the top-ranked agents for betweenness and eigenvector centrality in the three networks were central players in al-Muhajiroun, while 75 percent were false positives. In the final analysis, based on the improved thesaurus, nineteen out of the twenty-eight top-ranked agents for both measures were prominent activists and leaders in al-Muhajiroun. This represented 68 percent of the top-ranked agents in the final models, nearly a reversal of our preliminary results. Refining the thesaurus dramatically improved the accuracy of agent-based networks modeled by AutoMap, even though these networks were extracted from less than perfect data (see Tables 4.4 and 4.5).

Table 4.4. *Betweenness centrality in al-Muhajiroun using final thesaurus*

Rank	Network A	Betweenness Centrality	Network B	Betweenness Centrality	Network C	Betweenness Centrality
1	Omar Bakri	0.464	Omar Bakri	0.467	Anjem Choudary	0.305
2	Anjem Choudary	0.049	Anjem Choudary	0.148	Omar Bakri	0.170
3	Hassan Butt	0.024	Abdul Rahman Saleem	0.088	Abdul Rahman Saleem	0.094
4	Abu Izzadeen	0.017	Abu Izzadeen	0.042	Abu Izzadeen	0.065
5	Afzal Munir	0.010	Mohammed Babar	0.032	Abdul Muhid	0.052
6	Aftab Manzoor	0.008	Omar Khyam	0.023	Assad Ullah	0.027
7	Mohamad Omar	0.006	Abu Abbas	0.009	Abu Saalihah	0.022
8	Irfan Rasool	0.005	Khalid Kelly	0.008	Omar Khyam	0.020
9	Abdul Rahman Saleem	0.004	Omar Sharif	0.007	Afzal Munir	0.016
10	Asif Hanif	0.004	Abdul Muhid	0.004	Ishtiaq Alamgir	0.015
	Nodal False Positive		Relational False Positive		AM Activists/Associates	

Table 4.5. *Eigenvector centrality in al-Muhajiroun using final thesaurus*

Rank	Network A	Betweenness Centrality	Network B	Betweenness Centrality	Network C	Betweenness Centrality
1	Omar Bakri	0.662	Omar Bakri	0.654	Anjem Choudary	0.408
2	Anjem Choudary	0.377	Anjem Choudary	0.484	Abu Izzadeen	0.372
3	Hassan Butt	0.374	Abu Izzadeen	0.393	Omar Bakri	0.367
4	Afzeal Munir	0.316	Abdul Rahman Saleem	0.361	Abdul Rahman Saleem	0.345
5	Aftab Manzoor	0.311	Abdul Muhid	0.229	Abdul Muhid	0.335
6	Yasir Khan	0.297	Abu Uzair	0.227	Ali Behesti	0.323
7	Rubana Akhgar	0.276	Ishtiaq Alamgir	0.205	Ibrahim Hassan	0.318
8	Asif Hanif	0.250	Mizanur Rahman	0.200	Mizanur Rahman	0.318
9	Ibrahim Hasson	0.231	Abu Abdullah	0.196	Shah Jalal Hussain	0.304
10	Abu Maryam	0.222	Ibrahim Hassan	0.196	Omar Khyam	0.297
	Nodal False Positive		Relational False Positive		AM Activists/Associates	

VII. Conclusions

In an era of “big data,” automated coding applications like AutoMap have tremendous potential to increase our understanding of dark networks. Recognizing this potential, the Department of Defense and the intelligence community have invested substantial resources in these tools, allowing government analysts to rapidly mine large volumes of text searching for illicit nodes and the ties that bind them. In dark network analysis, where researchers face daunting challenges in gathering primary-source data, using such tools to harvest networks from news reports and other secondary data sources is particularly appealing.

As ethnographers engaged in computational network analysis of a dark network, we recognize the value of using AutoMap and ORA to extract and analyze networks of al-Muhajiroun activists from a large data set of news reports. While this chapter has emphasized the challenges we faced in using AutoMap, our research clearly benefited from it. Using AutoMap allowed us to process several thousand news articles, culling detailed models of al-Muhajiroun activists over three distinct periods. This would not have been possible using traditional hand coding methods. The amount of time and resources required for hand coding would have been enormous, greatly exceeding our capacity.

Although we spent considerable time – by our estimate 140 hours over seven weeks – creating and refining our concept list, the resulting thesaurus generated valid models of al-Muhajiroun activists that informed our thinking on the dark network. Particularly valuable were node and network-level centrality measures that helped us comprehend how al-Muhajiroun evolved in response to government enforcement efforts (Kenney et al., 2013; Kenney et al., 2015). This dynamic network analysis improved our understanding of al-Muhajiroun. As “thick” and nuanced as our ethnographic data were, the inferences we could draw from them were limited to the period we observed activists in London. Using AutoMap and ORA, our newspaper data could be partitioned and analyzed over a longer period. Other ethnographers would also benefit by using these tools to add a dynamic component to their analyses of dark networks.

Indeed, the moral of our story is not that ethnography is superior to computational network analysis, but that the two methods can be combined to enhance our understanding of dark networks like al-Muhajiroun. This speaks to the need for interdisciplinary research. Working together, computer and social scientists can leverage their skills in network analysis, statistical methods, and computer simulation, along with their in-depth knowledge of the case(s) under investigation. The strength of computational network analysis is that it

allows researchers to identify, extract, and measure dark networks from large volumes of data. The strength of ethnography is that it provides case-specific knowledge to help researchers determine whether the networks modeled are faithful representations of reality. While the Department of Defense quickly accepted the strengths of computational analysis, it has been slower to appreciate the value of ethnography. The Department's most prominent foray into anthropological research, the Human Terrain Teams, has been criticized by many professional anthropologists as being conducted on the fly by individuals with insufficient knowledge of local cultures and languages (Flintoff, 2010). Yet in the world of counterterrorism, where "kinetic" approaches to network disruption have become increasingly prominent, the moral and political implications of minimizing false positives are profound (Gerdes, 2014). Our research suggests that minimizing false positives in the analysis of covert networks requires complementing automated coding tools like AutoMap with traditional ethnography, however slow and cumbersome such methods are perceived to be.

Ethnography's strength is also a source of weakness, particularly when a single ethnographer, intentionally or otherwise, recreates his own mental map of the network by excluding agents he does not expect to see. Such confirmation bias may cause researchers to miss meaningful relationships in the network or to dismiss them as spurious. In our research we sought to control for this by "red flagging" individuals who raised questions and researching them more rather than dismissing them from the network at the outset. Many of these people turned out to be false positives, but more than once we uncovered individuals we did not realize were operationally associated with al-Muhajiroun. When this happened, we added them to the thesaurus. While much of what we learned from AutoMap confirmed our ethnographic understanding of the dark network, we also learned some new things as well, confirming our belief that computational network analysis could add value to our understanding of al-Muhajiroun.

Despite our best efforts, we cannot discount the possibility that we may have eliminated agents from the final thesaurus that should have remained, causing us to miss novel relationships. One step in addressing potential problems with confirmation bias is to formalize the red flag procedure we followed in our research by creating a separate thesaurus of flagged agents with the potential to falsify the ethnographer's mental model of the dark network. Agents in this thesaurus could be evaluated on a case-by-case basis. In accordance with the ontology of relationship types Gerdes offers in this volume's second chapter, meta-network analysis could serve to distinguish operational ties from ideological ties, a capacity that AutoMap and ORA possess.

A second step to address problems stemming from confirmation bias might be to include more ethnographers in the research design. One ethnographer with deep knowledge of the dark network could create the initial thesaurus, while another edits the final product. Alternatively, two ethnographers could work independently to create two thesauri that would be combined and vetted by a third expert. Variations on this theme are innumerable. Their feasibility, however, will depend in no small measure on DoD officials' willingness to accept ethnography's value for computational network analysis and to provide the resources necessary to fund such research. Dark network analysis works best when computer scientists and ethnographers work together to combine the strengths each brings to their research. These efforts may be slower than purely computational methods, but they are likely to produce more accurate models of dark networks.