MINING CRITICAL INFRASTRUCTURE INFORMATION FROM MUNICIPALITY DATA SETS: A KNOWLEDGE-DRIVEN APPROACH AND ITS APPLICATIONS

Wei-Ning Xiang
Department of Geography and Earth Sciences
University of North Carolina at Charlotte, NC 28223
wxiang@email.uncc.edu

William J. Tolone
Department of Software and Information Systems
University of North Carolina at Charlotte, NC 28223
wjtolone@uncc.edu

Anita Raja and David Wilson
Department of Software and Information Systems
University of North Carolina at Charlotte, NC 28223

Qianhong Tang, Ken McWilliams, and Robert K. McNally
Department of Geography and Earth Sciences
University of North Carolina at Charlotte, NC 28223

ABSTRACT

The practice of critical infrastructure protection (CIP) planning and management is often hindered by the lack of critical infrastructure (CI) information. This is due largely to the fact that an estimated 85 percent of all CI in the United States is in the private sector. Much of the information related to the security of critical infrastructures and essential to homeland security planning and emergency management is not customarily within the public domain—such information is commonly confidential, proprietary, and/or business sensitive. On the other hand, there is a great wealth of open-source data at the municipal level that is pertinent to critical infrastructures. These data are geographic and both collected and maintained by various departments in a municipal government to serve the general civil needs of the community. Examples include, but are not limited to, digital orthophotos, planimetric data sets, tax parcel data, and street centerlines. Because of their implicit connections to critical infrastructures, and owing to some of the important qualities they possess—high definition or precision (at a geographic scale of 1:2,400), refinement or completeness, and general availability—these data sets appear to be an attractive source of knowledge that await the practice of CI information mining. However, to date, no study has been reported that extracts only the most relevant CI information from this source, nor does a methodology exist that guides the practice of CI information mining on municipality data sets. A further challenge arises from the fact that the municipality data sets, as promising as they seem, are usually voluminous, heterogeneous, and even entrapping.

In this paper, we propose a knowledge-driven methodology that extracts CI information from open-source municipality data sets. Under this approach, pieces of deep yet usually tacit knowledge acquired from a group of CI domain experts are employed as keys to decipher the massive sets of municipality data and extract relevant CI information. More specifically, the knowledge worker first geographically renders the elicited knowledge on the generally available open-source information. The renderings are typically in the form of digital maps, attribute tables, and rule bases. The renderings are then reviewed by the expert(s) to make sure that the knowledge worker's interpretations of the human expert's knowledge agree with the expert's expectations. Upon the expert's feedbacks, adjustments are made. The corrected renderings are presented to the expert(s) for another round of validation. This interactive process continues until the expert(s) are satisfied with the results.

The proposed methodology was tested successfully on a municipality in the Southeastern United States. It is considered a viable choice for CIP professionals in their efforts to gather CI information for scenario composition and vulnerability assessment for several reasons. Firstly, it produces credible and useful results. Secondly, it is executable for most if not all municipalities across the United States as the required resources (data, computation, and knowledge acquisition, including domain experts) are generally available. Thirdly, it is adaptable and thus may be applied to any CI domain in any municipalities or regions in the United States. Lastly, methodology results can be readily refined and updated.

1. INTRODUCTION

A critical infrastructure (CI) is an array of assets and systems that, if disrupted, would threaten national security, economy, public health and safety, and way of life. These include, but are not limited to, utilities, medical facilities, public transportation, telecommunication networks, landmarks, buildings, and public spaces. In recent years, unfortunately, critical infrastructures have become symbolic targets as well as the mass casualty opportunities for terrorist attacks (Bolz *et al.* 2002, p.85). For instance, the World Trade Center is a symbol of America's capitalism and economic influence, the Pentagon is a symbol of America's military strength, and the railway station in Madrid represents a node in a geo-political network. Many critical infrastructures promote the congregation of people, which increases their attractiveness to terrorist acts. Because of the dual identity of critical infrastructures and the high level of vulnerability they bear, critical infrastructure protection (CIP) has topped the list of priorities in the practice of homeland security planning in the United States (Terner *et al*. 2004; Thieman, 2004). Since the tragic events of September 11th, 2001, CIP drills have become an integral part of every counter-terrorism exercise across the country (Thieman, 2004).

An essential task in CIP planning is the assessment of CI vulnerability with respect to the threat of potential terrorist attacks. For such a task, a set of *scenarios* is widely regarded in both academic and professional communities to be the best form for such assessments (Garrick 2002). Unlike predictions which project CI vulnerability with probability, a scenario set bounds the range of vulnerabilities by connecting initiating event, or initial conditions, to undesired end states (different levels of damage) with the sequence of events linking the two (Garrick 2002, p.421). Functionally, a scenario set is both a bridge that connects the process of CIP analysis and modeling with that of CIP planning, and a cognitive apparatus that stretches people's thinking and broadens their views in the practice of CIP. This dual function entitles a scenario set to be a favored member of a family of instruments for CI vulnerability assessment and CIP planning. It also explains the popularity of scenario sets in CIP drills and counter-terrorism exercises - they serve as a foundation for emergency response maneuvers (Thiemann, 2004).

Despite its recognized advantages, however, the use of scenarios in CIP and homeland security planning has been hindered by the difficulties in meeting its informational needs.

In composing scenarios, a scenarist requires certain categories of information, each corresponding to one of the five components of a scenario (Xiang and Clarke, 2003). These five components are: (1) alternatives—the range of potential actions, or the spectrum of incidents/events; (2) consequences—the immediate and cumulative effects (physical, ecological, economical, and social) that each alternative would have on an area's security, economy, public health and safety, and way of life; (3) causations—the causal bonds between alternatives and consequences; (4) timeframes—the periods of time between occurrence of the alternatives and the sequence of the consequences; (5) geographical footprints—the place-oriented blueprints of alternatives, and the anticipated marks of their ramifications on an area's geography. The last component—hardly unique—is so pivotal to both the composition and utilization of CIP scenarios that it indeed becomes a hallmark that distinguishes CIP scenarios from their counterparts in business, industry, and at times, the military.

Satisfying these informational requirements for composing CIP scenarios is by no means a light task. A major difficulty comes from the fact that an estimated 85 percent of all CI in the United States is in the private sector. Much of the information related to the security of critical infrastructures and essential to homeland security planning and emergency management is not customarily within the public domain - such information is commonly confidential, proprietary, and/or business sensitive (Terner *et al*. 2004). The Protected Critical Infrastructure Information (PCII) Program, launched by the U.S. Department of Homeland Security (DHS) in February 2004 under provisions of the Critical Infrastructure Information Act of 2002 (CII Act), enables the private sector to voluntarily submit infrastructure data to the federal government. The federal government assures that competitive data will be protected from public disclosure until and unless a PCII determination is made that the information does not meet PCII requirements.[1] Because of this voluntary nature, the success of this milestone development depends on collaborations from the private sector. Its efficacy remains to be seen.

On the other hand, there is a great wealth of open-source data at the municipal level that is pertinent to critical infrastructures. These data are geographic and both collected and maintained by various departments in a municipality government to serve the general civil needs of the community. Examples include, but are not limited to, digital orthophotos, planimetric data sets, tax parcel data, and street centerlines. Because of their implicit connections to critical infrastructures, and owing to some of the important qualities they possess—high definition or precision (at a geographic scale of 1:2,400), refinement or completeness, and general availability—these data sets appear to be an attractive source of knowledge that await the practice of CI information mining.[2] However, to date, no study has been reported that extracts only the most relevant CI information from this source, nor does a methodology exist that guides the practice of CI information mining on municipality data sets. A further challenge arises from the fact that the municipality data sets, as promising as they seem, are usually voluminous, heterogeneous, and even entrapping.

In this paper, we propose a knowledge driven methodology for CI information mining on municipality data sets. Under this approach, pieces of deep yet usually tacit knowledge are first acquired from a group of CI domain experts. These pieces of knowledge are then employed as keys to decipher the massive sets of municipality data and extract the CI information. The proposed methodology was tested successfully on a municipality in the Southeastern United States, and can potentially be applied to other municipalities and regions.

[1] http://www.dhs.gov/dhspublic/interapp/editorial/editorial_0404.xml.

[2] Data mining (also known as Knowledge Discovery in Databases—KDD) is defined as "The nontrivial extraction of implicit, previously unknown, and potentially useful information from data" (Frawley et al., 1992).

The remainder of the paper is organized as follows. Section 2 outlines the methodology. Section 3 details its implementation through a case study. Section 4 describes the outcomes of the case study. Section 5 concludes with the merits of the methodology, and suggests further research.

2. A KNOWLEDGE DRIVEN METHODOLOGY FOR CI INFORMATION MINING

The methodology proposed here gathers CI information through an intertwined process of knowledge acquisition, rendering, and validation (Figure 1). After an initial search, a knowledge worker, the person who solicits, interprets, and renders a human expert's knowledge, prepares a questionnaire and sends it to a human expert or, in some cases, a group of experts, for answers. In addition, the knowledge worker may schedule both structured and unstructured interviews with these expects. The knowledge worker then geographically renders the elicited knowledge on the generally available open-source information. The renderings are typically in the form of digital maps, attribute tables, and rule bases. The renderings are then reviewed by the expert(s) to make sure that the knowledge worker's interpretations of the human expert's knowledge agree with the expert's expectations, or at least, do not "violently contradict any strong held feelings" of the expert(s) (Keeney and Raiffa, 1976, p.271). Upon human expert's feedbacks, adjustments are made to refine the maps, attribute tables, and rule bases. The corrected renderings are presented to the human expert(s) for another round of validation. This interactive process continues until the expert(s) are satisfied with the results.

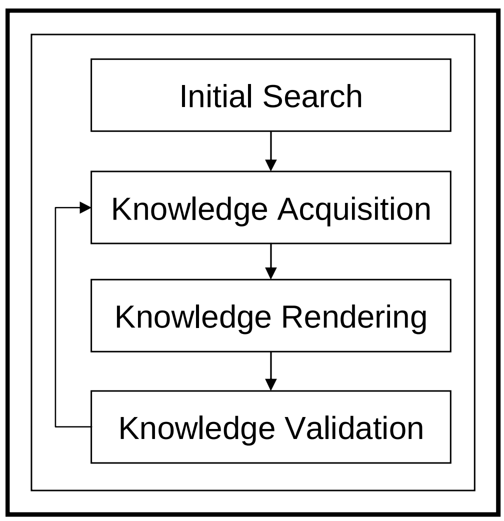

Figure 1. A methodology for critical infrastructure information mining

The remainder of the section provides a detailed account of this methodology. In some instances the initial search may generate sufficient knowledge such that a first rendering may occur prior to contacting domain experts (see Section 3.4). This initial rendering then serves as input to domain expert discussions.

2.1. Initial Search

The process of CI information mining begins with an initial search. This search involves a set of three tasks: i) a survey of all the publicly available information related to the domain; ii) the identification of the individual domain expert(s) for each CI component; and iii) the development of a list of questions for the interviews of the domain experts.

Publicly available information usually resides in open source locations such as generally accessible databases, journals, books, newspapers, the Internet, and even word-of-mouth. Much of the information from these sources is non-confidential, insensitive to business interests, and bearing little insights into the key issues related to CIP planning and management. However, both the information and the search process are necessary and valuable for CI information mining. They help knowledge workers become familiar with the problem domain; and they enable knowledge workers to identify individual domain experts and develop questionnaires in order to proceed to the next stage of the information mining process.

In should be noted that the thoroughness of the initial search contributes significantly to the success of the subsequent stage of knowledge acquisition. First of all, it makes the interviews, structured and unstructured, more productive while reducing the time commitment of domain experts, who usually have limited resources and in many cases must justify the hours spent with the knowledge workers. Secondly, domain experts are more likely to be forthcoming with procedural knowledge when knowledge workers are well prepared. This forthcoming nature of domain experts is independent of whether the knowledge worker utilizes a questionnaire during the interview.

2.2. Knowledge Acquisition

The task of knowledge acquisition deals with two levels of CI knowledge. Surface knowledge is the most common information offered when seeking information on a subject. Surface knowledge is also referred to as declarative knowledge (Turban and Aronson, 2001). Such knowledge includes facts and figures that can usually be found easily at the beginning of knowledge acquisition, and in some cases, during the stage of initial search. Examples include, but are not limited to, schematic diagrams of a CI system hierarchy, iconographical representations of CI roles, and functional relationships among objects within a CI layer. The CI proprietors usually are willing to give this generalized information to interested parties as it satisfies most inquiries and raises little corporate security concerns. Deeper procedural knowledge (Turban and Aronson, 2001), on the other hand, such as safety protocols, problem diagnostic procedures, and emergency response plans, are usually not disseminated to outside seekers for security concerns. Furthermore, the spatial element in both declarative and especially procedural knowledge, that is, information about the geographic footprints of CIs, either as individual layers or as components of a public utility system (a traffic control system, for example), is usually regarded as highly secure and business sensitive. Consequently, neither the data about locations and geographic networks of CI objects, nor the knowledge about proximity and spatial interactions among CI objects is readily available.

Various methods have been suggested for knowledge acquisition. These include, but are not limited to, the method of “familiar tasks” (Duda and Shortliffe, 1983; Mittal and Dym, 1985; Stefik et al, 1982), structured and unstructured interviews (Hoffman, 1987; Weiss and Kulikowski, 1984), limited information tasks (Hoffman, 1987), constrained processing tasks (Hoffman, 1984, 1987; Klein, 1987), and the method of “tough cases” (Doyle, 1984; Hoffman, 1987). Among these methods, interviews (both structured and unstructured) are the most common and widely accepted approach to knowledge elicitation (Hoffman, 1987; Millet and Harker, 1990; Xiang and Whitley, 1994). Although interviews are usually more time-consuming and less efficient than other methods (such as limited information tasks and constrained processing tasks), domain experts generally feel more comfortable with interviews and, therefore, are less hesitant to offer opinions (Ibid.). For example, during an interview conducted through our case study, the domain expert offered a detailed physical description of a critical node within the telecommunication network, providing details on square footage, parking area, number of stories and surrounding land use. This invaluable tacit knowledge would have remained undiscovered without such an interview.

At this stage of CI information mining, a combination of structured and unstructured interviews are used to acquire deep procedural knowledge, following an approach developed by Xiang and Whitley (1994, pp.283-288). The unstructured interview consists of an informal discussion with the domain experts, where the knowledge worker conducts an unscripted conversation and slowly builds his/her understanding of the infrastructure. Interviews with these individuals usually last about an hour and consist of an informal question and answer session. The structured, more formal interview is usually conducted afterwards, where the knowledge worker submits a list of detailed questions to solicit the domain expert’s explicit input. A structured interview often further refines the knowledge acquisition process, providing greater clarification to knowledge extracted during the informal interviews.

Ideally, a group of experts from each domain should be interviewed to minimize individual biases. Under this multi-expert approach, the knowledge acquisition process is usually conducted in either an interactive group setting or through individual interviews with domain experts. A group setting is considerably more time efficient because a knowledge worker can solicit input from multiple domain experts in a single session. However, there are several drawbacks associated with this method, including domination by one individual, lack of focus, and limited participation by all members (Moore 1987). With individual interviews, on the other hand, the knowledge worker will need to consolidate all of the information obtained from the domain experts. This process can be cumbersome, and time consuming - especially when the information is contradictory

(Roth and Wood 1990). In practice, therefore, the ideal multi-expert approach often yields to a sub-optimal single-expert approach for pragmatic reasons.

Another way to de-bias elicited knowledge is to use a shared information approach that encourages communications among the knowledge workers. During the interview process, each knowledge worker is usually assigned to a specific domain, working individually with a domain expert. Under this shared information approach, the knowledge workers share with one another the knowledge they solicited from various domain experts through informal discussions at regular intervals. These “cross-domain” discussions can provoke questions that could prove useful in follow-up interviews with the respective domain experts. In our case study, for instance, a final, formal presentation of knowledge acquisition results proved to be an extremely useful communication session for both knowledge workers and domain experts. The knowledge workers benefited from the additional informal knowledge validation provided by domain experts as well as information obtained from co-worker presentations. It also provided an excellent opportunity for domain experts of the various utility sectors to discuss their interdependencies.

2.3. Knowledge Rendering

Knowledge rendering is a process in which knowledge workers interpret the knowledge elicited from domain experts and express this understanding using an appropriate representation. More specifically, in the case of CI information mining, knowledge rendering refers to the exercise through which knowledge workers develop a *spatial representation* of the infrastructure layer based on the knowledge acquired from the domain expert(s). In our case study, for instance, a domain expert supplied the knowledge worker with a detailed description of the physical characteristics of a critical node in a telecommunication network (for example, parking lot size, building height, etc.). To render this piece of knowledge spatially, the knowledge worker first selected all the service provider’s properties from a publicly available tax parcel data file. The worker then superimposed these selected land parcels onto digital orthophotos of the study area, and removed all the parcels that do not match the physical characteristics. The worker finalized the rendering by saving all the remaining parcels on a new map. The map is referred to as a *proxy* or *surrogate* data layer in that it is not a map of the real data *per se*. Instead, it is a geographic rendering of the domain expert’s knowledge on the open-source municipal data (tax parcels and digital orthophotos, in this case) as interpreted by the knowledge worker. In addition to maps, the rendering of domain expert’s knowledge can take other forms, such as iconographic representations, flowcharts, and tables.

2.4. Knowledge Validation

During the final phase of the CI information mining process, the knowledge worker submits the newly rendered proxy data layer to the domain expert for validation. A series of structured or unstructured interviews are scheduled, where the domain expert is asked to verify that the spatial renderings are consistent with the knowledge he/she provided in the previous interview(s). The proxy data layers submitted to the domain experts are usually presented as digital maps. These layers are continuously updated and refined during the process based on feedback from the domain expert(s). This phase is complete when the domain experts are satisfied with the knowledge rendering presented and have verified that the knowledge worker did not misrepresented any of the knowledge.

3. IMPLEMENTATION OF THE METHODOLOGY: A CASE STUDY

The above-described methodology was initially developed to mine information from four critical infrastructures in a municipality in the Southeastern United States. These infrastructures include the electric power grid, telecommunication, natural gas, and transportation networks. As an illustration of the application of the methodology, this section provides a detailed description of mining exercise for information about the natural gas supply system. The outcomes are presented in Section 4.

A pipeline network supplies natural gas to residential and commercial customers throughout the metropolitan area. Commercial pipelines transmit large volumes of natural gas at a high capacity of 420 PSI (pounds per square inch, 1 PSI = 6895 Pascals).[3] Residential pipelines distribute small volumes at a capacity of 70 PSI or lower.

3.1. Initial Search

Because of the security concerns with gas pipeline data, it is not surprising that the initial search for publicly available data yielded a modest result. Among the most relevant are the pipeline system's general structure and a digital map of the major transmission pipelines from Texas to New Jersey. In addition, the knowledge worker collected information about the domain experts who are affiliated with the natural gas provider in the region, and the information on the land parcels designated as easement for natural gas providers.

3.2. Knowledge Acquisition

Through a series of structured and unstructured interviews with the domain expert, the knowledge worker obtained two important pieces of knowledge about the locations of the local pipeline network and the physical characteristics of the regulator stations—the nodes on the network. During a visit to the company, the knowledge worker was also given an image file, a screen shot from a GIS data file, of the local pipeline network. From the image file, it was clear that the easement areas on the property parcel data set are major pipelines in the municipality area.

3.3. Knowledge Rendering

The process of knowledge rendering consisted of three steps: construction of the regional pipelines (higher than 70 PSI); identification of the regulator stations, and creation of a local pipeline network (70 PSI and lower). This process proceeded under the guidance of the knowledge acquired from the domain expert.

[3] In figures and maps, PSI is referred to as "psig" --"pounds per square inch pipe".

In constructing the regional pipelines, the knowledge worker first superimposed a map layer of the easement areas on top of the digital orthophotos, and found some interesting physical characteristics of these easement areas: flat, open, linear, around 15 feet in width, either a little higher or lower than the surrounding ground surface. With the image file of the local pipeline network provided by the natural gas company as a reference, the knowledge worker then created a proxy of the regional pipelines by connecting those easement areas.

To identify the locations of the regulator stations, the knowledge worker first selected from a tax parcel database all the land parcels that are owned by the gas company and put them on a new map. The knowledge worker then superimposed the map on the digital orthophotos to examine, within each land parcel owned by the natural gas company, whether the physical characteristics of the regulator stations described by the domain expert were found. Typically, a regulator appears to be a small structure on a digital orthophoto with an electrical generator in the back. Another characteristic is that regulators are usually located at regular intervals along the pipeline. The knowledge worker then placed a node on the identified regulator to represent its location. Finally, all the nodes were saved on a new map layer.

The creation of local distribution pipeline network (70 PSI and lower) was guided by rules acquired from the domain expert. These rules are (1) 90% of the pipes follow the collector and neighborhood streets, and roughly 10% of the streets in a municipal street centerline file do not coincide with pipelines; (2) pipelines do not follow major highways—interstates and arterials; and (3) newly developed neighborhoods that do not have natural gas supplies are shown on the image file that the company provided as empty areas. Applying these rules to the street centerline files of the study area, the knowledge worker produced a map of the proxy distribution pipeline network.

3.4. Knowledge Validation

The completed renderings of the gas pipeline network and regulator stations proxy data layers were submitted to the domain expert for validation. Again, because of security concerns, the domain expert did not explicitly confirm whether the renderings were “right” or “wrong”. He instead directed the knowledge worker’s attention to places where erroneous assumptions in the pipeline placement had occurred. Based on this feedback, the knowledge worker made adjustments on the pipeline alignments on the digital maps, and checked on the tax parcel database to assure that the updates were thorough and complete.

It should be noted that during the process of knowledge acquisition, rendering, and validation, spatial techniques equipped in GIS and open-source data sets, especially, digital orthophotos played a significant role. By showing the geographic objects in their actual forms, rather than cartographic symbols (points, lines, areas, or pixels) [Lillesand and Kiefer 2000], digital orthophotos provide an advanced tool that was far more effective and efficient than the traditional maps and aerial photos for knowledge rendering and validation. The use of GIS enables the knowledge workers to manage the complex task of rendering an entire utility network. It helps track various features such as transmission poles, transmission stations and gas pipelines all within a same spatial domain. In this fully integrated system, all of the features are related by location and multiple data layers can be dynamically linked for visual analysis and for map production (Harder 1999).

4. PRODUCTS OF THE INFORMATION MINING PROCESS

The information gathered through the above processes falls into two interrelated categories. They are information about the functionality of a CI system and information about the geography of CI system structures. In this section, a description of the resulting information about the natural gas distribution system is given as an example of the products of the mining process.

Figure 2 shows the hierarchy of the natural gas pipeline network. Level 1 represents the main national gas pipeline (e.g., the West/East pipeline running from Texas to New Jersey). The 175 PSI lines at Level 2 transmit natural gas to the municipal area. The Level 3 network (70 PSI or lower) distributes natural gas to the local service areas.

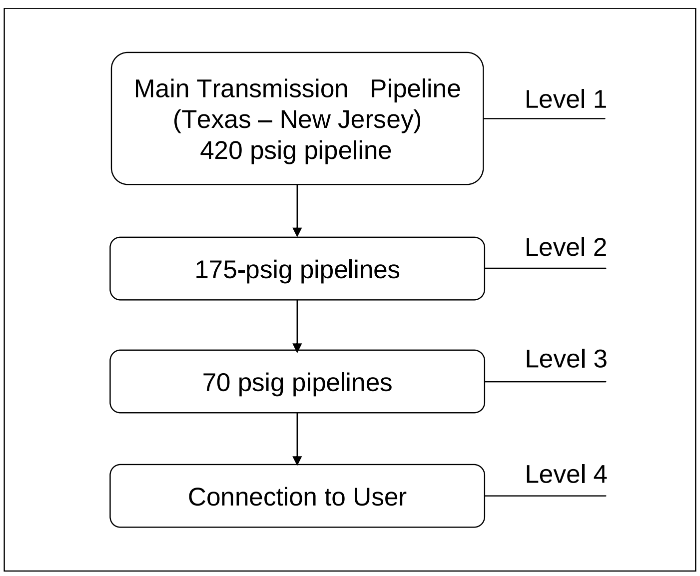

Figure 2. A functional hierarchy of the natural gas distribution system

A geographic representation of the natural gas distribution network in the study area is provided in Figure 3. The thickest line in the lower right corner of the map represents the main national gas pipeline. The regional pipeline (in blue) circles the municipal boundary and distributes gas—via regulator stations—to the local services pipelines (in dark green).

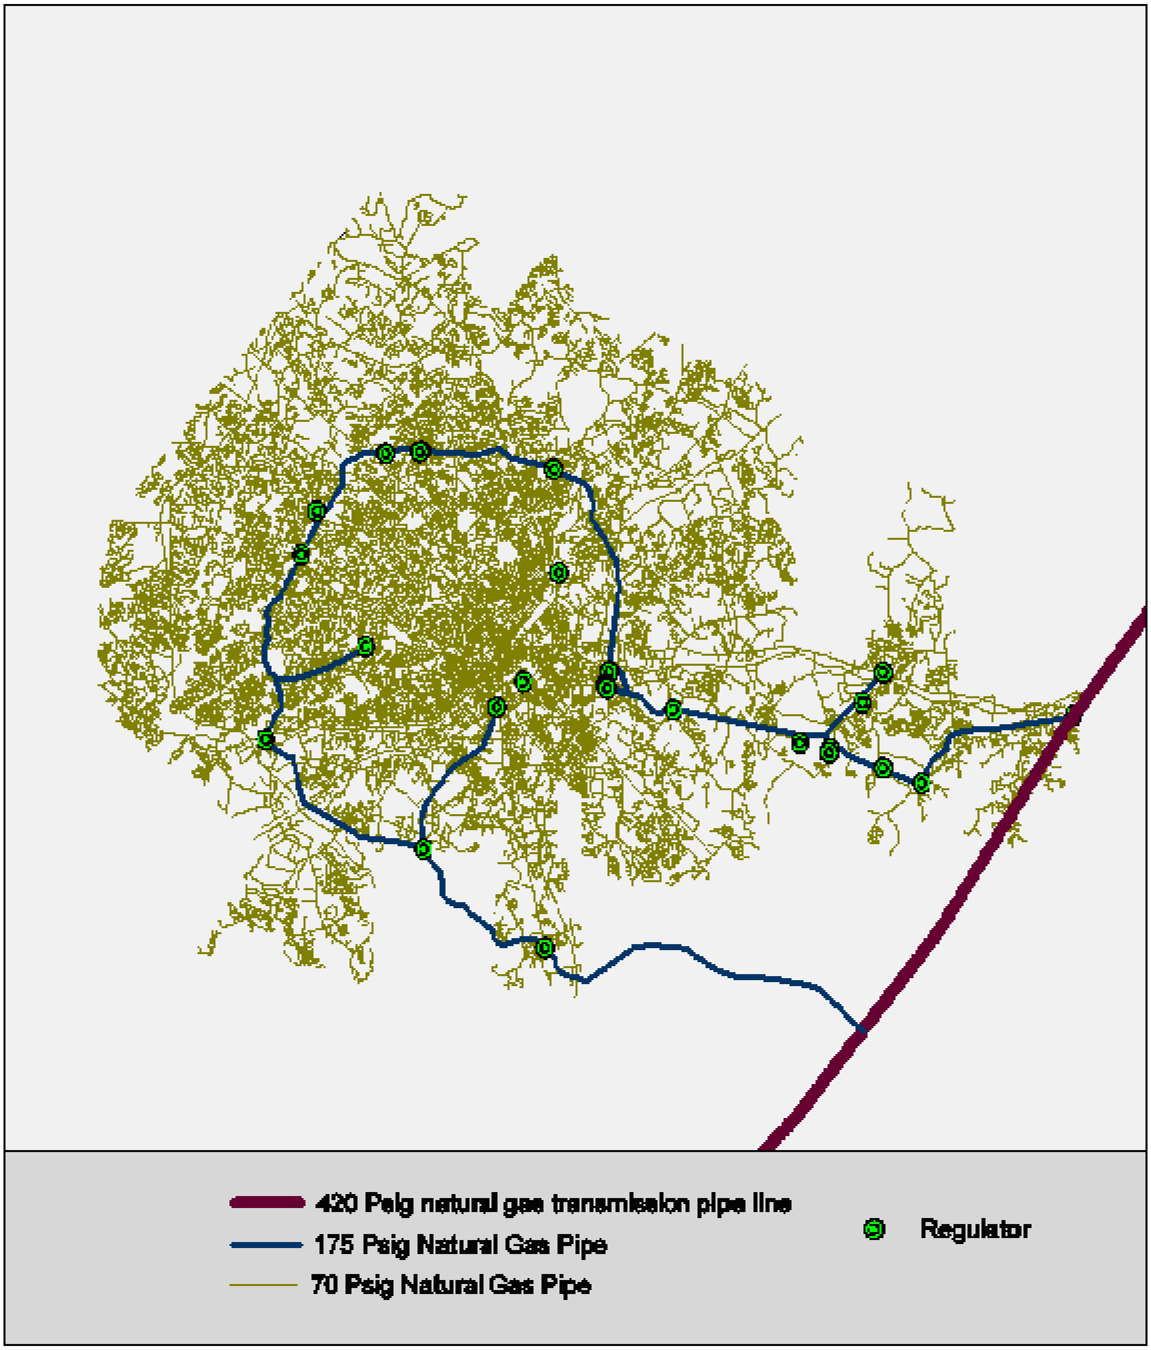

Figure 3. A Geographic Rendering of the Natural Gas Distribution System

5. CONCLUSIONS

We contend that the proposed methodology is a viable choice for CIP professionals in their efforts to gather CI information for scenario composition and vulnerability assessment for several reasons. First, the methodology produces credible and useful results, as demonstrated through the case study (the CI information gathered was credible in the eyes of domain experts). Second, the results are also useful for CIP scenarists in composing scenarios. The CI information from this case study has successfully supported a CI simulation project funded by a federal agency (Tolone 2004). Third, the methodology is executable for most, if not all, municipalities across the United States as the required resources for data, computation, and knowledge acquisition, including domain experts, are generally available. Fourth, the costs associated with methodology implementation are modest, most of which are related to the time and expenses of knowledge workers. Fifth, the methodology is adaptable and can be applied to any CI domain, although only four domains of CI information were involved in this case study. Finally, not only can methodology results be readily refined and updated, but they can also be incorporated into a real CI database should it become available under the Protected Critical Infrastructure Information (PCII) Program.

6. ACKNOWLEDGEMENTS

The authors are grateful to the following individuals at the University of North Carolina at Charlotte for their contributions to the project: Bei-Tseng Chu, Mirsad Hadzikadic, Vikram Sharma, Deepak Yavagal, Robin Gandhi, Katie Templeton, Gustavo Borel Menezes, Stuart Phelps, Jocelyn Young, Paul Smith, and Huili Hao.

7. REFERENCES

Benbasat, I., Dhaliwal, J.S., 1989. "A framework for the validation of knowledge acquisition" *Knowledge Acquisition*, 1, 215-233.

Bolz, F., Dudonis, K., Schulz, D., 2002. *The Counterterrorism Handbook: Tactics, Procedures, and Techniques.* (2nd edition). Boca Raton, FL: CRC Press.

Casazza, Jack and Delea, Frank. 2003. *"Understanding Electrical Power Systems"* John Wiley &Sons, NY

Doyle, J., 1984. "Expert Systems without Computers or Theory and Trust in Artificial Intelligence" *AI Magazine*, 5(2), 59-63.

Duda, R.O. and Short life, E.H., 1983. "Expert Systems Research". *Science*. 220, 261-276.

Frawley, W., Piatetsky-Shapiro, G., and Matheus, C., 1992. "Knowledge discovery in databases: an overview. *AI Magazine*, 13, 213-228.

Garrick, J., 2002. "Perspectives on the Use of Risk Assessment to Address Terrorism". *Risk Analysis*, Vol. 22, No. 3. 421 – 423.

Harder, Christian, 1999. *Enterprise GIS for Energy Companies.* ESRI Press Inc.

Hoffman, R.R., 1984. "Methodological Preliminaries to the Development of an Expert System for Aerial Photo Interpretation (TRETL-0342)". Fort Belvoir, VA: Engineer Topographic Laboratories.

Hoffman, R.R., 1987. "The Problem of Extracting the Knowledge of Experts from the Perspective of Experimental Psychology". *AI Applications in Natural Resources Management*. 1(2), 35-48.

Keeney, R.L. and Raiffa, H., 1976. *Decision with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley & Sons.

Klein, G.A., 1987. "Applications of Analogic Reasoning". *Metaphor and Symbolic Activity*, 2, 201-218.

Lillesand, T., Kiefer R., 2000. *Remote Sensing and Image Interpretation.* New York, NY: John Wiley and Sons, Inc.

Millet, J. and Harker, P.T., 1990. "Globally Effective Questioning in the Analytic Hierarchy Process". *European Journal of Operational Research*. 48, 88-97.

Mittal, S. and Dym, C.L., 1985. “Knowledge Acquisition from Multiple Experts”. *AI Magazine*. 6(2), 32-36.

Moore, C.M. 1987. *Group Techniques for Idea Building*. Newbury Park, CA: Sage Publications Inc.

Roth, M and Wood, C., 1990. “A Delphi Approach to Knowledge Acquisition from Single and Multiple Experts”. *Trends and Directions in Expert System.*

Stefik, M., Aikins, A., Balzer, R., Benoit, J., Birnbaum, L., Hayes-Roth, F., Sacerdoti, E., 1982. “The Organization of Expert Systems: A tutorial” *Artificial Intelligence*. 18, 135-173.

Terner, M., Sutton, R., Hebert, B., Bailey, J., Gilbert, H., and Jacqz, C., 2004. Protecting Critical Infrastructure. GeoIntelligence, March 1, 2004 (http://www.geointelmag.com/geointelligence/article/articleDetail.jsp?id=90043)

Thieman, K.J., 2004. *The Use of Scenarios and Geographic Information Systems and Technology (GIS &T) in Counter-Terrorism Exercises and Planning*. Master’s Thesis, Department of Geography and Earth Sciences, University of North Carolina at Charlotte, Charlotte, North Carolina (copies available by contacting the first author of this paper).

Tolone, W.J., Wilson, D. , Raja, A., Xiang, W.N., Hao, H., Phelps, S. and Johnson, E.W. 2004 “Critical infrastructure Integration Modeling and Simulation.” In *Proc. of the 2nd Symposium on Intelligence and Security Informatics (ISI 2004). Lecture Notes in Computer Science #3073.* Springer-Verlag.

Turban, E., Aronson, J.E., 2001. *Decision Support Systems and Intelligent Systems*. Upper Saddle River, NJ: Prentice Hall.

Weiss, S., and Kulikowski, C., 1984. *A Practical Guide to Designing Expert Systems*. Totowa, NJ: Rowman and Allanheld.

Xiang, W.-N. and Whitley, D.L., 1994. “Weighting Land Suitability Factors by the PLUS Method”. *Environment and Planning B: Planning and Design.* 21, 273-304.

Xiang, W.-N. and Clarke, K.C., 2003. “The Use of Scenarios in Land Use Planning”. *Environment and Planning B: Planning and Design.* 30, 885-909.