

A Generalized Query Framework for Geospatial Reasoning

Martin Michalowski

Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292 USA
martinm@isi.edu
Advisor: Dr. Craig A. Knoblock

Problem Statement

The ability to reason over geospatial entities using publicly available information is greatly enhanced by the abundance of data sources on the Internet. Traditional data sources such as satellite imagery, maps, gazetteers and vector data have long been used in *geographic information systems* (GIS). However, incorporating non-traditional sources such as phone books and property tax sites brings to light integration issues not previously considered. The ability to answer different types of queries about geospatial entities further complicates the problem. For example, assume we are provided with an image and consider the query “What is the name of this street?”. There exist many different ways of answering this query depending on the types of data sources and operations available to a system. Three possible plans which answer this query are shown in Figures 1, 2, and 3 .

For each plan, we are provided with the latitude and longitude of the center of the image and a point identifying which street we are interested in. The first plan would query a vector data source¹ which provides a street name(s). The second plan queries a gazetteer² which returns a set of named points. These points could be used to query an online phone book data source for a street name(s). Finally, the third plan retrieves an image from an image data source, overlays this image on the map, and extracts the characters corresponding to a street name(s).

These plans show there exist components that can: produce and pass new data to subsequent components in the plan, produce an incomplete set of data, or produce no data at all. Therefore, when constructing an execution plan, a system needs to consider not only the operations available to it but also the data produced during execution. There is no guarantee that a valid execution plan will produce a valid query answer. The notions of *contingency planning* or *dynamic plan refinement* need to be incorporated into the planning system.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹Given an area of interest, a vector data source provides meta-data and positional information about geospatial entities.

²A gazetteer is a data source that produces points with types and location coordinates when provided with an area of interest. See <http://www.alexandria.ucsb.edu> for an example.




Figure 1: An execution plan that uses a vector data source to provide street name(s)




Figure 2: An execution plan that uses a gazetteer and an online phonebook to provide street name(s)

Traditionally, a mediator system such as the one proposed by (Thakkar, Ambite, & Knoblock 2005) could answer this query. However, in answering the query mentioned above, the traditional mediator would produce one plan of execution and assert that a valid plan produces an answer to the query. Yet, if during the execution of the plan in Figure 2 the *Gazetteer* operation produces no data, we should have the recourse to abandon that plan and begin executing either plans one or three. This capability is not available in current mediated systems. This research attempts to incorporate this type of reasoning into a generalized query-answering framework.

Research Plan

The goal of this research is to provide a framework that can integrate traditional data sources (satellite imagery, vector data) with non-traditional data (phone books, information web sites). This framework provides a query-answering capability when supplied with a query relating to geospatial entities within a specified area. Using a model of the available operations and sources, the framework dynamically generates an execution plan based on the type of query, the preconditions of operators, and the input bindings on sources. Planning requires reasoning both at the data and the operations level. This can be accomplished by interleaving traditional planning techniques with data integration paradigms used for data-level reasoning.

By intertwining data level and operator level reasoning, execution plans can be refined dynamically depending on the data seen up to that point. For example, a traditional planning system is provided with a domain model which




Figure 3: An execution plan that overlays an image on a map then extracts the characters to provide street name(s)

describes all of the available operations. This description includes the preconditions and effects associated with each operator and a planning system is very good at composing plans by using this information. On the other hand, a plan composed by reasoning at a data level is more robust and flexible, allowing for dealing with operator failures or data inconsistencies. Merging the operator level reasoning capabilities provided by traditional planning techniques with a data level reasoning capability provides the tools to develop a better underlying query answering framework.

The proposed query-answering framework uses a planning approach. It represents the user query as a goal (or a set of subgoals) and uses backtracking or query reformulation to come up with a plan that answers that user query. By interleaving planning and execution, data produced by the operators is validated and a plan is refined or re-planned as necessary.

Additionally, a query answer is improved or augmented when possible. Specifically, once a user query is answered, there exists additional information relevant to the query created from the sum of the operations used in the plan. This additional information would be used to augment the query answer. It should be noted that inconsistencies across data from different data sources, multiple possible answers, and other important issues make this a difficult and intriguing problem.

Progress to Date

A novel approach to solving the integration issues seen in labeling buildings in satellite imagery was proposed in (Michalowski & Knoblock 2005). This approach leverages techniques developed in the planning community by casting the problem as a constraint satisfaction problem (CSP). It also presents a unique way of using both explicit and implicit information in publicly available data sources. Finally, it showed that successfully integrating traditional and non-traditional data sources can be done using a planning framework. The work will be integrated as an operation available to the generalized framework proposed in this research and current efforts are focused on improving its flexibility for solving different types of labeling problems and its scalability to larger problems.

The main challenges remaining in the proposed research are in the area of query representation and operator modeling. Current efforts concentrate on establishing a representation that is general enough to answer many different types of queries but specific enough to allow for comprehensive reasoning to be done at the data level. A prototypical system used an abstract representation of the data similar to the approach in (Knoblock 1995) and the FF planner (Hoffmann & Nebel 2001) as the planning backbone for the system. Intuition suggests that a purely abstract representation may not

be conducive to the goals of the system. For example, having identified three of four buildings correctly for a given query, an abstract representation makes it difficult to re-plan for the remaining unidentified building. A hybrid representation, one that combines abstract (set) and specific (tuple level) views of the data is the ideal approach. The development of such a representation is the main focus of current research efforts.

Determining the viability of using existing methods to include the reasoning capability described in this paper is also an important step in the research. Methods introduced in integrating planning and execution (Golden 1998), knowledge-based planning (Petrick & Bacchus 2004), and information gathering (Kwok & Weld 1996) provide valuable insights into ways of solving the research problem. Furthermore, extending the Prometheus mediator (Thakkar, Ambite, & Knoblock 2005) to support the desired reasoning capability is an interesting avenue also being explored.

The analytical goals mentioned above determine if any existing systems can both: (1) represent user queries as goals and (2) be used to implement a hybrid planning and execution module. The prototypical system showed that if need be, various different representation, planning, and execution approaches can be combined in interesting and novel ways.

References

- Golden, K. 1998. Leap before you look: Information gathering in the PUCINI planner. In *Artificial Intelligence Planning Systems*, 70–77.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- Knoblock, C. A. 1995. Planning, executing, sensing, and replanning for information gathering. In Mellish, C., ed., *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1686–1693. San Francisco: Morgan Kaufmann.
- Kwok, C. T., and Weld, D. S. 1996. Planning to gather information. In *13th AAAI National Conf. on Artificial Intelligence*, 32–39. Portland, Oregon: AAAI / MIT Press.
- Michalowski, M., and Knoblock, C. A. 2005. A constraint satisfaction approach to geospatial reasoning. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*.
- Petrick, R. P. A., and Bacchus, F. 2004. Extending the knowledge-based approach to planning with incomplete information and sensing. In *International Conference on Automated Planning and Scheduling*, 2–11.
- Thakkar, S.; Ambite, J. L.; and Knoblock, C. A. 2005. Composing, optimizing, and executing plans for bioinformatics web services. *The VLDB Journal, Special Issue on Data Management, Analysis, and Mining for the Life Sciences* 14(3):330–353.