Using a Geographic Information System for Qualitative Spatial Reasoning about Trafficability

CPT James J. Donlon[1]   Kenneth D. Forbus

Qualitative Reasoning Group
Northwestern University
1890 Maple Avenue
Evanston, IL, 60201, USA
donlonj@csl.carlisle.army.mil,forbus@ils.nwu.edu;

Abstract

Modern geographic information systems (GIS) are powerful quantitative tools for geospatial reasoning. We show how a GIS can be used as a metric diagram to support qualitative spatial reasoning. We illustrate the technique with the problem of reasoning about trafficability -- an entity's capability for movement through some terrain. Many common geospatial reasoning tasks rely heavily upon estimates of trafficability, and in some domains, off-road trafficability is of primary concern. Maps rarely provide the appropriate representation of space for such problems. We show how GIS technology can be usedas a metric diagram from which place vocabularies appropriate to the task can be computed. We describe the domain theories that enable the computing of these place vocabularies, and provide models for qualitative trafficability reasoning based on these descriptions. The results suggest that the power of qualitative representations and reasoning can be brought to a variety of geospatial applications, riding on the progress made by the commercial world in GIS software.

Introduction

Modern geographic information systems (GIS) provide powerful and useful facilities for digitizing and mapping representations of geographic space. They are widely used for performing complex transformations of this data, and for solving common geospatial problems, (e.g., placement of structures under constraints). However, most GIS computations are quantitative, relying on visualization tools and extensive user interactions to provide the qualitative insights that users need. In contrast, most human reasoning about geographic space appears to use a qualitative interpretation of that space. We believe that qualitative spatial representations can be useful for reasoning about geospatial problems. The metric diagram/place vocabulary (MD/PV) model uses a combination of quantitative and qualitative representations, decomposing space in task-specific ways into the kind of regions that are intuitive to human users and useful for qualitative reasoning [4,6]. Our key observation is that GISs can be used as metric diagrams. That is, they can directly answer quantitative queries (which is their standard use), and they can be used to construct place vocabularies, i.e., task-dependent qualitative spatial representations, which is a novel use for such systems.

In this paper we use the task of geospatial reasoning for trafficability analysis problems to show how GISs can be used in qualitative spatial reasoning. We begin with a brief description of GIS systems. Then we describe how the GIS digital representations of terrain are combined with symbolic reasoning to automatically produce qualitative descriptions of terrain. The nature of trafficability problems is discussed next, followed by a description of our domain theory for solving such problems. We outline the successful results obtained with our implemented system, and propose directions for future work.

[1] Current address: Knowledge Engineering Group, United States Army War College, Center for Strategic Leadership, 650 Wright Avenue, Carlisle, PA, 17013, USA

Brief review: How GIS systems work

GIS systems enable the production and analysis of digital representations for geospatial problems. They accept geographic data from many sources and in many formats, and provide two- and three-dimensional visualizations of this data from a multitude of perspectives. They are used in many problems, such as mapping, land management, city planning, and vehicle routing. Although variations exist, in general GIS systems encode their geospatial data and accompanying attribute information in a relational database.

GIS systems organize this information in *layers*. Each layer contains geospatial information of some particular type with a consistent set of spatial entities (points, arcs, or polygons) that cover the region of interest. For example, it is not unusual for a digital terrain data set to consist of polygon coverages that describe vegetated areas, water bodies, etc., arc coverages for roads, rivers, sewer lines, etc., and point coverages that represent facilities, nodes in transportation networks, etc.

The properties used to distinguish regions within a layer are either symbolic labels (e.g., vegetation-cropland vs. vegetation-woodland), numerical values, or numerical ranges. What differentiates a region in a GIS layer as distinct is that all of the attribute values within this region are uniform.

GIS as metric diagram

In the MD/PV model, the *metric diagram* plays the role that diagrams do for human spatial reasoning: they provide concrete descriptions from which we can easily answer spatial questions. Humans routinely decompose continuous space into discrete regions according to criteria relevant to the desired reasoning. Everyday geographic reasoning provides many examples of this, e.g., zones for urban development, districts for government. *Place vocabularies* provide this kind of qualitative representation. The computation of place vocabularies from metric diagrams models the human use of visual processing to extract/impose qualitative spatial distinctions on diagrams or scenes, and the ability to use combinations of purely qualitative and metric information in spatial reasoning.

The digital data and geometric processing provided by GIS technology means they can be used directly as a metric diagram for geospatial reasoning. The mixture of symbolic and numeric information in a GIS is similar in kind to the representations found in earlier metric diagrams (e.g., [4,6]), but optimized for geospatial problems. The library of geometric and database routines provided with high-end systems are more than adequate for implementing routines to handle concrete, quantitative spatial questions, thus satisfying the first constraint on metric diagrams.

The second service of metric diagrams, the creation of place vocabularies, is partially provided by the GIS ability to produce new layers that quantize space into polygons and arcs. The questions that must be solved to apply a GIS to a qualitative geospatial domain are

1. What place vocabulary(ies) are needed for the task?
2. How can they be computed from the available data?

The rest of this paper shows how we have answered these questions for an important class of geospatial reasoning, *trafficability problems.*

The problem: Trafficability

Trafficability is a measure of the capability for vehicular movement through some region. It is a relationship between some entity (capable of movement) and the area through which it moves. Whether an area is trafficable for the entity, and the measure of how trafficable the area is for the entity, depend on the interaction between the entity's mobility characteristics and the relevant terrain attributes. A simple example is planning a car trip, where one plots routes through a network of roads and reasons about the relative merits of different routes (e.g., total travel time, possible effects of bad weather, etc.). A more complex example is plotting a route for a multi-vehicle safari through open country. While our methods are applicable to either on-road or off-road movements, we are particularly concerned with trafficability reasoning for the more difficult case of cross-country movement.

Some simple on-road trafficability problems are best solved using known quantitative methods. For example, route finding and travel time problems in a fully specified road network can be easily solved with built-in GIS algorithms for path finding and flow optimization. However, many important trafficability problems do not lend themselves to such solutions. Taking into account weather-related factors or partial knowledge about the terrain make even on-road trafficability problems difficult using standard solution methods. Off-road, cross-country trafficability problems are far more complex, and there are no off-the-shelf tools that automatically solve them. Figuring out if an area can be reached with a particular vehicle and exploring migration patterns are two examples of tasks requiring off-road trafficability reasoning.

The particular focus of our work is to military trafficability problems -- specifically, the trafficability of military vehicles for cross-country movement. Military estimates rely heavily on an understanding of the trafficability of terrain. In the military domain, qualitative trafficability determinations about regions of space are combined to inform a variety of estimates, such as identifying potential areas through which large forces might move, or determining where a force might perform specific types of operations.

The basis for most military terrain analysis is map data and terrain analysis surveys. Use of maps as the sole basis for trafficability estimates is problematic because of the amount and variation of representations presented on a single topographic map. It is also very difficult and labor-intensive to integrate data gathered in field surveys with the topographic map's representation of the terrain. To manage this information overload, the military's terrain analysis processes rely heavily on diagrams that identify places that have uniform significance or characteristics relevant to the desired analysis. These diagrams (*overlays*, in the military vernacular) are produced by outlining the regions that are uniform for the relevant characteristics on transparent sheets, registered to the underlying map. The resulting overlays are then used extensively for producing further estimates. They also provide a representation for communicating about the area of operations.

Two particular products of military terrain analysis inspired the place vocabularies we computed. These are the *complex factor overlay* (CFO)[8], and the *combined obstacle overlay* (COO)[9].

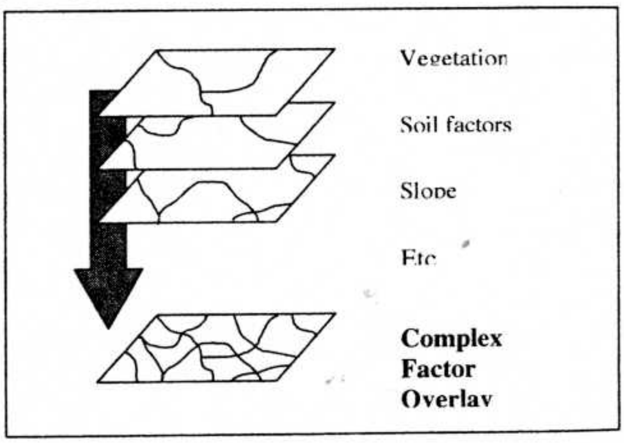

Figure 1. Complex Factor Overlay

A CFO is the cumulative partitioning of space according to criteria from several terrain analysis categories. It is derived by creating overlays for each category (i.e., slope, soil factors, vegetation, surface materials, and hydrology), and then combining them to identify the areas with homogenous characteristics across all overlays. The result is a partitioning of space that is as detailed as the available terrain data permits (Figure 1). This overlay allows the analyst to compute estimates of vehicle speeds for each region, based on vehicle capabilities and formulae that relate those characteristics to the terrain in the region.

In contrast, the COO is a characterization of the terrain's effects on vehicular movement, characterized in a much more general way. Terrain in each region is said to be *unrestricted* (U), *restricted* (R), or *severely restricted* (SR) for movement in the relevant context, e.g., movement of a brigade equipped with tracked and wheeled vehicles. (Water regions are identified separately.) Thus, the production of a COO results in fewer individuated places, as adjacent U, R, and SR regions are aggregated (Figure 2).

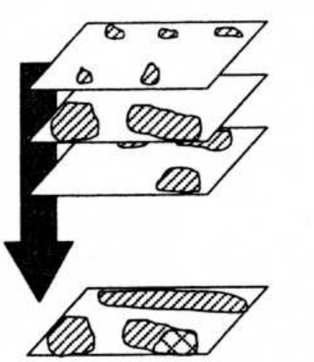

Figure 2. Combined Obstacle Overlay

The COO is primarily used for determining where military units are likely to be moving and conducting operations. The echelon of the unit under consideration helps to determine what size areas of R and SR terrain are significant enough to be included on the overlay.

Both the CFO and the COO are examples of place vocabularies made explicit in human practice. They are qualitative decompositions of space used to help solve particular classes of tasks. The qualitative reasoning challenge is to automate the production of these descriptions.

Using a GIS to construct place vocabularies

Places are individuated by where some important property or combination of properties is uniform. Each polygon in a GIS layer is defined to be uniform with respect to whatever property (or properties) is associated with that layer. In a sense the GIS already implements a place vocabulary for any digital terrain data representation, albeit of an extreme form.

The key idea in using a GIS for qualitative spatial reasoning is to (a) figure out how to compute the properties needed for a task from the available data and (b) derive a new layer that decomposes the terrain in terms of uniformity in those properties. The new layer constitutes a place vocabulary for that characterization of the terrain. For example, in the combined obstacle place vocabulary, the characterization is whether a piece of terrain is U, R, or SR, and those regions where this property is uniform correspond exactly to the regions found in a Combined Obstacle Overlay.

To compute the properties required, we developed a domain theory for trafficability, expressed in CML [1] augmented with KIF [10]. To create the new GIS layers corresponding to the COO and CFO, we developed algorithms that used the geometric processing library of ARC/INFO [15] to do the necessary computations. We first describe the domain theories, and then we describe the GIS procedures.

Trafficability domain theories

Reasoning about vehicle trafficability requires knowledge of terrain and its effects on movement, vehicles and their capabilities to move on and off road, and how terrain affects the movements of these vehicles over the terrain. We developed domain theories to address each of these three areas, as well as domain theories that represent knowledge about how the digital representation of the terrain is encoded. We will briefly describe each of these theories, and provide some representative definitions and axioms as illustration. The size of these domain theories is indicated in Table 1.

Domain Theory	# Predicates	# Axioms	# facts
Trafficability	87	164	19
Terrain	504	731	16
Military Vehicles	285	388	0
GIS	45	151	25
Totals:	921	1434	60

Table 1: Size of domain theories

Trafficability

Trafficability is a binary relationship between the terrain and the vehicle. This relationship can be analyzed and represented either qualitatively or quantitatively, depending on the task.

Qualitative expressions of trafficability are expressions of whether a vehicle is capable of effectively traveling through a region, and the qualitative effects the terrain will have on that movement. For example, the determination that a region is trafficable for a movement of some kind (represented here as a movement episode [2]) relies on determining that no part of the path is untrafficable.[1] The path of the movement must provide sufficient clearance for the physical extent of the movement, and a suitable surface for the movement of the vehicle:

```
(defRelation trafficable (?P ?M) :=
 (and (movement-episode ?M)
   (path ?P)
   (not (or (weather-denied ?P ?M)
     (exists (?S ?T ?V)
       (and (segment-of ?S ?P)
            (terrain-of ?S ?T)
            (vehicle-of ?M ?V)
            (or (insuff-clearance ?P ?V)
                (untraff-surface ?T ?V)
  )))))))
```

[1] On the "negative polarity" of considering a segment trafficable: It is more natural and convenient to specify which conditions prevent movement, and assume their non-existence unless there is evidence to the contrary. For example, in the absence of proof of weather-denied, we should assume it is false.

A region's surface is considered unsuitable if the slope, vegetation, soil conditions, or surface materials of the region are prohibitive, or if the region is a body of open water.[2]

```
(defRelation untrafficable-surface  (?T ?V)
  := (or (too-steep ?T ?V)
         (too-vegetated ?T ?V)
         (weak-soil ?T)
         (obstructed-surface ?T)
         (dvalue (obstacle-category ?t) W)))
```

Further expanding the illustration, a region is too vegetated for the cross-country movement if the vegetation is both too closely spaced for the vehicle to drive around and too thick for the vehicle to override:

```
(defRelation too-vegetated (?T ?V)
  := (and (vegetated ?T)
          (and (< (stem-spacing ?T)
                  (min-turning-radius ?V))
               (> (stem-diameter ?T)
                  (override-diameter ?V)))))
```

These axioms can be very useful for identifying critical relationships that constrain possible movements.

We also use compositional modeling techniques [3,5] to express qualitatively the effect the terrain in a region will have on a movement through it, such as the effects of vegetation:

```
(defmodelfragment
 movement-rate-vegetation-effects
 :participants ((M :type movement-episode)
                (P :type path))
 :conditions ((possible-trajectory-of M P))
 :consequences
  ((qprop+ (movement-rate M P) (SS P))
   (qprop- (movement-rate M P) (SD P))))
```

Where SS represents the distance between stems or trunks in a vegetated area, and SD represents the average diameter of those stems. This axiom gives us the basis for making qualitative judgements such as, all other things being equal, a

[2] We use the relation `dvalue` to specify the value of a discrete variable. Similarly, we use `nvalue` to specify the value of a continuous parameter. This allows us to disentangle the semantics of these relations from =, `equal`, `eql`, etc. for reasoning.

vehicle can travel faster through scrub than it can through dense woods. This is just the kind of common-sense judgement that human analysts routinely apply when interpreting descriptions of terrain and making predictions about movements over that terrain.

The effects of weather also play a significant role in trafficability analysis. We can use similar models to account for these effects, enabling the domain theory to take advantage of reasoning about weather information, if available:

```
(defmodelfragment
 trafficability-weather-effects
 :participants ((M :type movement-episode)
                (P :type path))
 :conditions ((possible-trajectory-of M P))
 :consequences
  ((qprop- (movement-rate M P)
           (visibility P))
   (qprop- (movement-rate M P)
           (wetness P))
   (qprop- (movement-rate M P)
           (accumulation-of rain M))
   (qprop- (movement-rate M P)
           (accumulation-of snow M))
   ...))
```

For problems that demand more quantitative trafficability determinations, we draw from the U. S. Army's cross-country movement analysis technique for estimating vehicle speed over terrain regions.[8] In practice, these places are derived from the CFO overlay described above. Vehicle speed is calculated for a given vehicle for each type of region based on a series of mathematical transformations. Cross-country movement speed (CCM) is calculated based on the vehicle's top speed, unconstrained, which then is degraded based on the cumulative effects of slope (F1) and surface configuration (F2), vegetation (F3), soil effects (F4), and surface materials (F5):

```
(deffunction ccm-speed-dry
             (?veh ?geo-area) :=
 (* (F1/2 ?veh ?geo-area)
    (F3 ?veh ?geo-area)
    (F4D ?veh ?geo-area)
    (F5 ?veh ?geo-area)))
```

The variable `?geo-area` is bound to a GIS polygon. Determining each of the above factors involves a number of steps. For example, determining the vegetation factor (F3) involves the combination or comparison of six other factors. One of these factors is V1a, a factor that accounts for a vehicle's ability to drive around the trees in a vegetated area:

```
(deffunction Vla (?veh ?geo-area) :=
  (cond ((<= (Vla-calc ?veh ?geo-area) 0) 0)
        ((<= (Vla-calc ?veh ?geo-area) 1) 1)
        (t (Vla-calc ?veh ?geo-area))))

(deffunction Vla-calc (?veh ?geo-area) :=
 (/ (- (SS ?geo-area)
       (SD ?geo-area)
       (width ?veh))
    (- (min-turning-radius ?veh)
       (* 4 (width ?veh)))))
```

Although the reasoning involved in this estimate is largely quantitative, qualitative descriptions of space allow the reasoner to perform these calculations when the terrain data is missing or incomplete. Normally this calculation is not possible if a value is missing for any one of the properties that are involved in calculating the factors. For example, while our GIS data set contained the structure for the full range of military terrain attributes needed to determine these factors (e.g., stem spacing (SS) and stem diameter (SD), above) the coverage rarely contained values for these attributes. To fill in missing information, our reasoner uses the named characterization of the region, e.g. EB020 (scrub/brush), to conduct default reasoning about the missing quantities.

Finally, we formalized the characterization of terrain as used in the combined obstacle overlay. Recall that in the COO terrain areas are identified as U, R, SR. The following are some of the axioms that define severely restricted terrain:

```
(defrelation severely-restricted-terrain
             (?t ?v)
  := (and (dvalue (obstacle-category ?t) SR)
          (transportation-device--vehicle
            ?v)))

(forall (?t)
 (=> (and (landform ?t)
          (or (>= (percent-slope ?t) 45)
              (and (< (SS ?t)
                      (meters 6))
                   (>= (SD ?t)
                      (meters .15)))))
  (dvalue (obstacle-category ?t) SR)))
```

```
(forall (?t)
 (=> (and (waterbody ?t)
          (or (> (water-current ?t)
                 (feet-per-second 5))
              (> (water-depth ?t)
                 (feet 4))
              (> (military-gap-width ?t)
                 (meters 17))))
  (dvalue (obstacle-category ?t) SR)))

(forall (?x)
  (=> (and (terrain-feature ?x)
       (or (dvalue (vegetation-type ?x) K)
           (dvalue (vegetation-type ?x) J)
           (dvalue (vegetation-type ?x) I)))
  (dvalue (obstacle-category ?t) SR)))

(forall (?x)
 (=> (weak-soil ?x)
  (dvalue (obstacle-category ?t) SR)))
```

Terrain

The trafficability domain theory relies on a detailed representation of terrain, including:

Physiography - Representations of land formations and their presentation (e.g., hills, valleys, and dunes). Characteristics used to describe landforms include: Percent of slope, slope intercept frequency, elevation, major types of landforms.

Hydrography -- Water and drainage associated features (e.g., river, stream, lake, sabkhat). Characteristics include: Water depth, water current, banks

Vegetation – Presence of plant life (i.e., forest, swamp, grassland), with characteristics including: vegetative roughness, stem spacing, stem diameter, vegetation type, canopy closure, tree-height

Surface Materials/Soils - Soil type and classification, and a classification of surface roughness types, including: Area soil type, primary soil partition, secondary soil characteristic, rating cone index, surface roughness

We have represented terrain in sufficient detail to conduct the kind of trafficability reasoning described above. This involves a general characterization of terrain features:

```
(defrelation vegetated (?x)
 :=> (and (geographic-area ?x)
          (<= (VR ?x) 1)
          (>= (SS ?x) 0)
          (>= (SD ?x) 0)))
```

as well as descriptions of terrain regions that have some uniformity of attributes:

```
(defentity scrub
  :subclass-of (vegetated-terrain-feature)
  :consequences ((uncultivated :self)
                 (sapling-growth :self)
                 (sparsely-vegetated :self)
                 (default-dvalue
                    (vegetation-type :self)
                      B2)
                 (vr-scrub :self)))
```

it involves default reasoning for missing information, e.g.,

```
(defrelation VR-scrub (?x)
  :=> (default-nvalue (VR ?x) .80))
```

and relevance to military trafficability analysis:

```
(defrelation
     severely-restricted-vegetation (?x)
  := (and (<= (SS ?x) 6)
          (>= (SD ?x) .15)))
```

Military vehicles

The vehicle theory contains an ontology of military vehicles, with an emphasis on their quantities and attributes relevant to cross-country mobility. Vehicle characteristics provided by the domain theory include maximum road speed, width, length, height, weight, on-road gradability, off road gradability, fuel capacity, fuel consumption (idle), fuel consumption (secondary), fuel consumption (cross-country), override diameter, vehicle cone index (1 pass), vehicle cone index (50 passes), minimum turning radius, load class, maximum gap to traverse, ground clearance, maximum step, maximum tilt, maximum gradient, and maximum straddle.

We account for a wide range of military vehicle types, which allows us to reason about a variety of trafficability scenarios. The domain theory ranges from representations of general classes of vehicle:

```
(defentity vehicle-wheeled
 :subclass-of
  (transportation-device--vehicle)
 :quantities
  ((number-of-wheels))
 :consequences
  ((default-dvalue (drive-type :self)
      wheeled)))
```

to specific vehicle types:

```
(defentity M-2
  ;; Bradley Infantry Fighting vehicle (BFV)
  :subclass-of
   (InfantryFightingVehicle
    MilitaryVehicle-Tracked
    MilVeh-amphibious)
  :consequences
   ((default-nvalue
     (max-road-speed :self) 66)
    (default-nvalue
     (fuel-Capacity :self) 175)
    (default-nvalue
     (fuel-Consumption-idle :self) 6.4

   ... Etc.
   )
  :documentation "The standard IFV of U.S. Army
mechanized infantry units.
Equipped with ...")
```

GIS Topology and Feature Coding

A theory of GIS topology is necessary to provide a way to describe and reason about the representation of the terrain. This gives us the vocabulary and definitions needed to produce the place vocabulary layers in the GIS.

For example, the domain theory contains knowledge about the way GIS manages data:

```
(defEntity gis-coverage-poly
 ;;; GIS coverage with polygon toplogy
 :subclass-of (gis-coverage))
```

It includes the geometric representations used to individuate regions of space:

```
(defEntity gis-polygon
 :subclass-of (gis-feature)
 :quantities ((gis-poly-area)
              (gis-poly-perimeter)))
```

And it includes axioms that define the relations between these regions, which allow us to express the qualitative spatial descriptions that make up the place vocabulary:

```
(forall (?poly1 ?poly2 ?arc)
 (=>  (and (gis-arc ?arc)
           (gis-polygon ?poly1)
           (gis-polygon ?poly2)
           (gis-right-of-arc ?arc ?poly1)
           (gis-left-of-arc ?arc ?poly2))
   (and (gis-adjacent-polygons
          ?poly1 ?poly2)
        (gis-adjacent-polygons
          ?poly2 ?poly1))))
```

Another GIS domain theory module provides a translation out of the format peculiar to the digital terrain data set, and into the common predicate vocabulary provided by the terrain and trafficability domain theories. It is a body of rules that translates feature ID codes such as EB020 to the appropriate predicate vocabulary, `(scrub GIS-POLY-021)`:[1]

```
(forall (?poly)
 (=> (gis-poly-vt ?poly EB020)
   (scrub ?poly)))
```

Because this translation is handled in a knowledge-based way, GIS data coded to some other format could be used just as effectively by providing the appropriate feature coding domain theory.

GIS Place Vocabulary Algorithms

Place vocabularies are computed in response to a request from a reasoning system. For the sake of context, it will be helpful to understand how we utilized a client-server relationship between the reasoning system and the GIS to enable this. The process begins when the reasoner has recognized the need for a place vocabulary (i.e., a COO or CFO). The request is sent to the GIS via remote procedure call. The GIS server executes scripts that perform the requested operation, produces the resulting place vocabulary in a format that the reasoning system can read, and signals the completion of the task to the reasoner. The reasoner can then access the place vocabulary through direct communication with the GIS (itself a DBMS) or by directly reading uncompressed data files. The reasoner thus has access to all of the information it needs.

[1] Our translation is for terrain data that is coded to the NIMA Vector Smart Map data Level 1 format (VMAP-1).[13]

We next outline the algorithms used for computing CFOs and COOs. The actual GIS code is written in Arc Macro Language (AML) for ARC/INFO version 7.x [15]; We summarize them in terms of the following operations:

- *Clip* creates a new coverage by copying the subset of an input coverage that overlaps a given rectangle. Polygons that intersect the sides of the rectangle are clipped to match the boundaries.
- *Union* creates a new coverage with polygons whose boundaries are the union of the polygons in the input coverages.
- *Dissolve* merges adjacent polygons within a coverage that have identical values for a property provided as input.

Algorithm for generating a Complex Factor Overlay

Since the CFO consists of the finest-grained distinctions that can be made based on the input data, computing it is relatively straightforward. Given a set of input coverages Ci, representing the factors and a rectangle representing the area of interest (AI),

1. Let CFO = a new empty coverage
2. For each input coverage Ci,
   2.1 Let C = Clip(Ci,AI)
   2.2 CFO = Union(CFO,C)

Given data that has already been pre-processed to meet military standards, this algorithm correctly computes a CFO by definition.

Algorithm for generating a Combined Obstacle Overlay

Computing the COO is slightly more complex because we must compute the appropriate level of abstraction. In addition to a set of input coverages, Cis, and an area of interest AI, a number corresponding to the echelon forwhom the analysis is being done (e.g., brigade, division, corps, etc.) must be provided:

1. Let COO = a new empty coverage
2. For each input coverage Ci,
   2.1 Let Ci = Clip(Ci,AI)
   2.2 For each polygon P in Ci, assign obstacle-category W, U, R, or SR according to the constraints of the domain theory.
   2.3 Let Ci = Dissolve(Ci, obstacle-category)
3. For each input coverage Ci, let COO = Union(COO, Ci)
4. For each polygon P in COO, let obstacle-category(P) be the most restrictive of the values for obstacle-category of the corresponding polygons in the input coverages.
5. Let COO = Dissolve(COO, obstacle-category)
6. Remove all polygons in COO that are too small for the echelon under consideration

Essentially, this algorithm finds the maximal regions that are severely restricted, restricted, and water, and then prunes out regions that are too small for the given echelon.

Empirical Results

We used these techniques to generate CFOs and COOs to answer terrain analysis and trafficability questions relevant to planning and conducting military operations. The GIS data set we used represented the terrain in the Straits of Hormuz region. The relevant coverages in this terrain data included vegetation, hydrology, slope, and road network. This array of coverages, and the (predominantly desert) terrain in the area provided an opportunity to test these techniques, and the early results are promising. Inspection by military personnel of the CFOs and COOs produced was used to judge correctness, and plausibility of trafficability results judged in terms of producing results consistent with U.S. Army practice. In all areas tested, correct CFOs and COOs were created, and trafficability questions (e.g., maximum speed in particular regions) produced correct answers.

Producing CFOs and COOs gave us the opportunity to produce place vocabularies that correspond to authentic qualitative descriptions produced by Army analysts, and to model the well-established terrain analysis practices that use those descriptions. The payoff is the automaton of this practice, authentic results conducting terrain analysis and trafficability tasks, and compelling explanations grounded in qualitative reasoning. While military analysts currently spend hours or days producing the same overlays for an area of operations, this technique allows the same descriptions to be produced immediately, on demand. They can be produced in response to a human analyst, as well as in response to an automated reasoning process.

Discussion

We have shown that a GIS can be used to produce representations for qualitative spatial reasoning. The geometric processing facilities of the GIS provide the capabilities in a metric diagram. Layers of a GIS computed with respect to task-relevant properties constitute place vocabularies. Thus the power of qualitative representations and reasoning can be brought to geospatial applications, riding on the progress made by the commercial world in GIS software.

There are several directions that should be explored next. First, further testing is needed, especially using terrain data that differs substantially from the arid Hormuz area. Our Hormuz data also lacked information about soils, surface configuration, and surface roughness, so we have not validated our techniques with those aspects of terrain data. Second, closely related problems for which trafficability plays an important role (e.g., analyzing the suitability of an area for an operation, accessibility, and path finding) should be tackled, to provide a more encompassing set of tests. Third, while we experimented with GIS scripts that would call our reasoning system while they were running, this was both inefficient and impractical, given the concurrent development of the domain theories, reasoning system, and GIS procedures Consequently, we hand-coded the knowledge from the domain theory into the GIS scripts. This is a sensible solution, given that doctrine changes very slowly, and the same scripts can thus be widely used. However, either on-line integrated reasoning or automatic production of GIS scripts may be worth exploring for other GIS tasks. Fourth, the domain theories and place vocabularies should be expanded to include other spatial representations that are related to trafficability that would be a great value to terrain analysts. These would demand additional geometric processing, and different models in order to derive them. The avenue of approach overlay is an example of such a task, and a natural extension of the work we have done here.

There are many other geospatial tasks where using a GIS for qualitative spatial reasoning is likely to provide important benefits. Examples include anthropology (e.g. the migrations of cultures) and agriculture (e.g. crop selection). Regardless of the application, harnessing the geospatial computing power of commercial GIS should contribute significantly to creating qualitative spatial reasoners.

Acknowledgements

This research was supported by the Defense Advanced Research Projects Agency, under the High Performance Knowledge Bases program, and by the Air Force Office of Sponsored Research. We thank Tom Mostek, Rob Rasch, and Jeff Usher for their help and useful discussions.

References

1. B. Falkenhainer, A. Farquhar, D. Bobrow, R. Fikes, K. Forbus, T. Gruber, Y. Iwasaki, & B. Kuipers. CML: A Compositional Modeling Language. Knowledge Systems Laboratory, KSL-94-16, September 1994.
2. Cohn, A., & Randall, Z.(1992). An Interval Logic for Space Based on 'Connection'. Proceedings of the 10th European Conference on Artificial Intelligence. 394- 398.
3. Falkenhainer, B. and Forbus, K. "Compositional Modeling: Finding the Right Model for the Job", Artificial Intelligence, **51** (1-3), October, 1991.
4. Forbus, K. "Spatial and Qualitative Aspects of Reasoning about Motion" Proceedings of the first annual conference of the American Association for Artificial Intelligence, August 1980.
5. Forbus, K. Qualitative Process Theory. *Artificial Intelligence*, **(24)**, 1984, pp 85-168
6. Forbus, K., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The CLOCK Project", *Artificial Intelligence*, **51** (1-3), October, 1991.
7. Forbus, K. 1995. Qualitative Spatial Reasoning: Framework and Frontiers. In Glasgow, J., Narayanan, N., and Chandrasekaran, B. *Diagrammatic Reasoning: Cognitive and Computational Perspectives.* MIT Press, pp. 183-202.
8. FM 5-33 Terrain Analysis, Headquarters, Department of the Army, July 1990.
9. FM 34-130 Intelligence Preparation of the Battlefield, Headquarters, Department of the Army, July 8,1994.
10. M. R. Genesereth & R. E. Fikes. Knowledge Interchange Format, Version 3.0 Reference Manual. Knowledge Systems Laboratory, KSL-92-86, June 1992.
11. Glasgow, J., Karan, B., & Narayanan, N. (Eds.). (1995). *Diagrammatic Reasoning.* Cambridge Mass.: AAAI Press/The MIT Press.
12. Hayes, P. (1979). The naive physics manifesto. In D. Michie (Ed.), Expert Systems in the Micro-Electronic Age Edinburgh University Press.
13. Military specification MIL-PER-89040, Vector Product Interim Terrain Data (VITD), May 8, 1995, http://164.214.2.59/publications/specs/printed/VITD/vitd.html.
14. Student Handout, "Tactical Commander's Development Course Battle Book", U.S. Army Command and General Staff College, Fort Leavenworth, KS
15. ARC/INFO Geographic Information System for Windows NT, version 7.2, ©Environmental Systems Research Institute, Inc.