

# Qualitative Spatial Interpretation of Course-of-Action Diagrams

Ronald W. Ferguson\*  
ferguson@cs.nwu.edu

Robert A. Rasch, Jr.†  
raschr@leavenworth.army.mil

William Turmel\*  
turmel@ils.nwu.edu

Kenneth D. Forbus\*  
forbus@ils.nwu.edu

\*Qualitative Reasoning Group  
Department of Computer Science  
Northwestern University  
Evanston, IL 60201, USA

†Battle Command Battle Lab (BCBL)  
415 Sherman  
Ft. Leavenworth, Kansas 66027, USA

## Abstract

This paper demonstrates qualitative spatial reasoning techniques in a real-world diagrammatic reasoning task: Course-of-Action (COA) diagrams. COA diagrams are military planning diagrams that depict unit movements and tasks in a given region. COA diagrams are a useful test bed for researching diagram understanding due to their composable symbology, their intrinsically spatial task, and their use across many types of military planning. We constructed two COA diagram interpreters using our qualitative spatial reasoning engine, GeoRep. The first system uses GeoRep to interpret individual COA glyphs. The second system, building upon the first, takes pre-classified symbol input and then uses GeoRep to describe geographic relationships implied by the symbol arrangements. This latter system, in a recent DARPA initiative, answered dozens of geographic queries about many different COA diagrams. This research shows that qualitative spatial reasoning, through tools like GeoRep, provides a useful substrate for complex diagrammatic reasoning.

## Introduction

A key characteristic of diagrams is their capacity to quickly communicate many spatial relationships. Even the simplest diagrams, such as drawn maps, convey a large number of relations through the size, orientation, and placement of visual elements. For this reason, we argue, diagrammatic reasoning lends itself to an approach based on qualitative spatial reasoning.

We present a case study of how qualitative spatial reasoning was used in one class of real-world diagrams. These diagrams, Course-of-Action (COA) diagrams, are used for military planning in many different situations and at many different echelon levels. A current military priority for COA diagrams is to integrate information from them with other information and reasoning sources—such as libraries of previously-created COA diagrams or military expert systems.

Our approach to interpreting COA diagrams emphasizes qualitative spatial reasoning performed using our spatial representation engine, GeoRep. The spatial nature of the task and the compossibility of the COA symbology make GeoRep's qualitative spatial reasoning a useful substrate for reasoning about COAs.

In the following section, we describe the nature of the COA domain, and show how qualitative spatial reasoning is useful in this domain. After describing GeoRep and how it works, we then describe GeoRep's role in two COA interpreters: one prototype that describes the contents of simplified COA diagrams, and a second system which performs geographic reasoning on full COA diagrams. This second system was used as a qualitative geographic reasoner in the DARPA High Performance Knowledge Bases initiative.

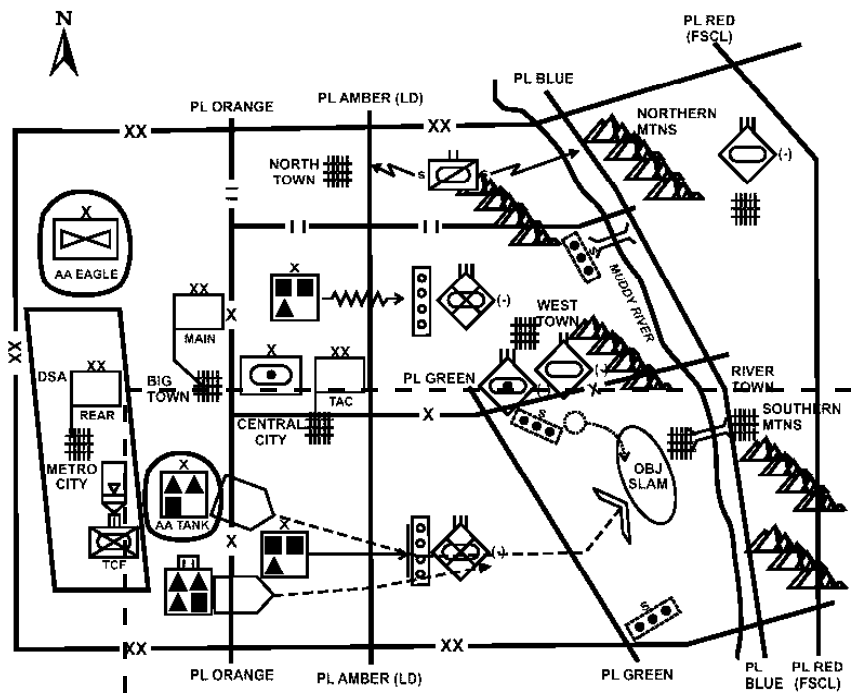


Figure 1: A course of action diagram (Department of the Army, 1997, Figure 5-5). For brevity, our description focuses on the dashed rectangle area, depicting the main attack (and two supporting attacks) on Objective SLAM.

Ferguson, R. W., Rasch, R. A., Turmel, W., & Forbus, K. D. (2000). Qualitative spatial interpretation of Course-of-Action diagrams, *Proceedings of the 14th International Workshop on Qualitative Reasoning*. Morelia, Mexico.

## Course of Action Diagrams

Course-of-Action (COA) diagrams provide an ideal testing ground for research into diagram understanding. COA diagrams have both an inherently spatial task and a broad and extensible visual symbology.

COA diagrams are military planning diagrams which relate a set of units and tasks to a geographic region. A COA diagram, given a rough map of a region's geographic features, depicts that region's military units and shows their assigned tasks (Figure 1). Along with units and tasks, the COA diagram depicts movement types (via arrow types and polyline symbols), available routes, topographical features, and other military tasks, such as blocking enemy movement. The diagram is accompanied by a written description of the intent and desired end state (Note: collectively, this text and the diagram constitute the whole "Course of Action"—we use the term "COA diagram" to refer to the diagram alone).

COA diagrams are interpreted qualitatively, especially when used in military planning, due to the need to continually adapt the plan as alternatives are considered. While COA diagrams may be drawn using modern drafting techniques, in many cases they are hand-sketched with grease pencils on large acetate sheets (sometimes overlaying a map). COA diagrams may be redrawn several times to remove irrelevant details, change the description level, or illustrate alternatives. In these cases, while exact measures are sometimes used (e.g., to estimate travel times), the diagram is mostly used to capture a set of qualitative spatial relationships in a way that allows quick assessment and modification.

The COA symbology is simply-drawn and broadly composable. Each visual symbol (or *glyph*) in a COA diagram (Department of Defense, 1999) can be captured in a few pen strokes and easily classified by its visual structure. Figure 2 shows standard COA glyphs for a boundary, a task organized unit, a minefield, a friendly and an enemy armor battalion, and a main attack on an objective area. Each glyph uses composable subparts. For an armor battalion (Figure 2(d)), the rectangle indicates a friendly unit, and the contained ellipse indicates an armored unit (it could also be a diamond (e), indicating an enemy unit). The two "antennae" above the unit indicate a battalion. One antenna indicates a company, an X indicates an entire brigade, and an XX an entire division. These same echelon markers transfer their meaning to other symbols in analogous ways, also indicating the echelon of the task force (b), the enemy force (e) and the border (a).

The concision and composability of the COA diagram symbology evolved from its long use as a visual vocabulary. Forms of the symbology were used for high-level military planning as early as the 1740's (Luvaas, 1966), and in the lower ranks since World War I. COA diagrams are now used by generals on down to company commanders. Time and technology have expanded the standard symbology, which now encompasses hundreds of glyph types defining myriad units, tasks, obstacles and boundaries.

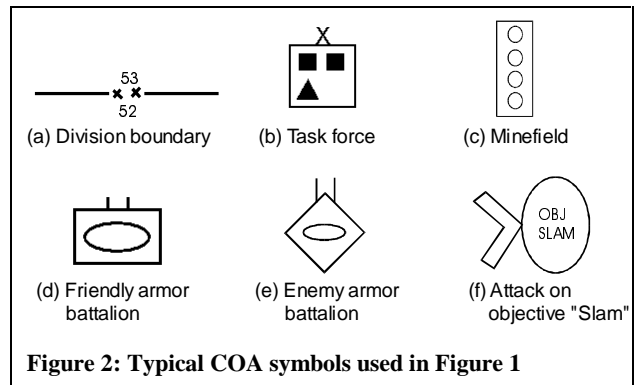


Figure 2: Typical COA symbols used in Figure 1

However, the symbology captures only part of COA diagrams' expressiveness. COA diagrams also communicate meaning via relative symbol placement. For example, Figure 1 depicts (within the dashed rectangle) three task force units attacking objective SLAM. Each unit is placed along an attack arrow, which assigns that unit to the main or supporting attack, and indicates the movement path. The paths cross an enemy minefield, indicating an enemy regiment holding position behind that minefield. Phase lines intersecting those paths (such as phase lines ORANGE and AMBER) show phases of planned movement during that attack (as described in the COA's statement). The diagram boundaries and their markings determine the level of the COA diagram (e.g., the outer boundary markings in Figure 1 indicate a division-level plan), and divide the region into "areas of operation" to which units are assigned responsibility.

In general, interpreting the meaning of COA diagrams requires two kinds of knowledge, both of which have qualitative spatial characteristics. Locally, it requires an understanding of COA glyphs and how they are composed. More globally, the reasoner must understand the implications of particular spatial relationships between glyphs, glyph placement relative to paths and boundaries, implicitly-defined directions of movement (the enemy unit "hiding" behind the minefield), and containment and adjacency relations between regions.

## GeoRep: A qualitative spatial reasoner

The simplicity and composability of COA diagrams give them their human utility, and also make COA diagrams good candidates for interpretation by a qualitative spatial reasoner. Thus, we constructed systems on top of our qualitative spatial representation engine, called GeoRep (Ferguson & Forbus, 2000).

GeoRep is a system for building diagrammatic reasoners using a low-level qualitative spatial representation as a substrate. As input, GeoRep takes a line drawing, given as a set of primitive visual elements in a vector graphics file. From this drawing, GeoRep creates a qualitative spatial representation of the visual relations found in the drawing. The representation is in terms of a place vocabulary defined through a visual domain theory.

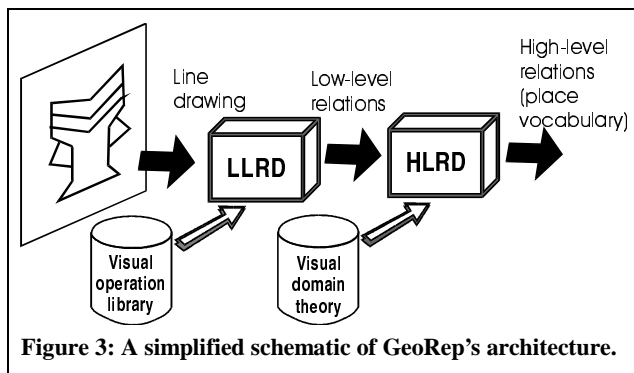


Figure 3: A simplified schematic of GeoRep's architecture.

GeoRep's architecture is shown in Figure 3. GeoRep's architecture contains two stages, the low-level relational describer (LLRD) and the high-level relational describer (HLRD). The LLRD handles the domain-independent representation of the line drawing. It detects and represents a large set of useful visual relations, including proximate elements, parallel line segments, polygons and polylines, connection relations, and containment relations. In general, these visual relations are those detected early in perception by universal visual routines (Ullman, 1984).

The HLRD, in turn, uses domain-specific rules that extend the LLRD's representation. These extensions include new visual relations and ways to recognize depicted domain symbols. The HLRD's output is a set of qualitative spatial relations that correspond to a specific task or type of analysis. For example, representation levels may include the LLRD's basic visual representation, more complex visual relations, a representation of the depicted items, or potentially even reasoning within the problem domain.

GeoRep's architecture is described in more detail in (Ferguson & Forbus, 2000).

### Applying GeoRep to the COA domain

We used GeoRep to build two reasoners in this domain.

Our first COA diagram interpreter uses GeoRep to recognize a subset of the COA symbology. This system explores how the composability of glyphs in COA diagrams is expressed through element shapes and qualitative spatial characteristics, such as containment. This reasoner, the *COA diagram describer* (COADD), works directly from simplified COA diagrams (Figure 4) drawn using the JavaFIG drawing program (Hendrich, 1999).

COADD uses a simplified subset of the COA symbology. Its domain includes basic unit and attack types, symbols for assembly, engagement, and objective areas, and minimal boundary lines. Because we were interested in the nature of the symbology and its composability, we did not work from bitmaps, but from vector graphics files (as is done with other GeoRep-based systems). While these diagrams were much simpler than typical COA diagrams, they still captured significant compositionality and expressiveness. For example, Figure 4 shows a dual attack by three brigade-level task forces on three objective areas.

COADD has a visual domain theory for COA diagrams written as HLRD rules. These rules recognize COA




Figure 4: Example from a Course-of-Action diagram. Three friendly brigade-level task forces attack three enemy positions. The main attack is against objective Buford, and the supporting attack is against objective Grant.

symbols using low-level structural relations. For example, friendly armor units are recognized as rectangles containing horizontally-oriented ellipses. These rules follow the compositionality of the COA symbology. For example, to take advantage of how echelon markers apply across many different symbol types, COADD's rules first find echelon markers (i.e., small groups of crosses or hash marks), and then use this information to assign echelon levels to specific units and boundaries.

Other rules in the visual domain theory infer the intent of units from other spatial characteristics, such as units' proximity to attack arrows. For example, when a unit is within an attack arrow, that unit is assigned to that attack task. Proximity to an assembly area determines the units gathered at that area.

A useful result for qualitative spatial reasoning is how COADD benefits from the LLRD's extensive low-level spatial description. Because the low-level spatial relations correspond to easily perceived and described visual relations, it provides a generative vocabulary for describing symbologies. The LLRD's spatial description thus allows the visual domain theory to be simple. COADD's visual domain theory contains 37 geometric rules to cover 18 basic object types and relations. The LLRD also allowed the system to be built quickly: an initial version of COADD, handling everything except a few difficult recognition tasks (recognizing heterogeneous arrow types and broken borders) was done in less than 10 person-days.

COADD works reasonably well, producing representations of units and tasks rich enough to infer the COA's general traits (though not to infer intent). Figure 5 contains a partial COADD representation for Figure 4.

We tested COADD's representational capability by using it to build a retriever for COA diagrams. This has an important potential application: when planners evaluate a new COA plan, examining similar old plans provides insight into the new plan's potential side-effects.

<u>Enemy unit inside area "Grant"</u>	<u>Subordinate units for the task force</u>	<u>Other</u>
(armor-unit unit-6)	(subordinate-maneuver-unit unit-10 unit-17)	(completed-minefield minefield320))
(enemy-unit unit-6)	(mechanized-infantry-unit unit-10)	<u>Representation links</u>
(military-unit unit-6)	(battalion unit-10)	<u>Bottom arrow</u>
(motorized-rifle-unit unit-6)	(friendly-unit unit-10)	(represents (composite <polyline:4> (axis-of-advance-support attack4))
(unit-at unit-6 (objective-area area19 "Grant"))	(subordinate-maneuver-unit unit-13 unit-17)	<u>Echelon marker</u>
	(mechanized-infantry-unit unit-13)	(represents (x-mark <segment:71><segment:72> (marker marker317 brigade))
<u>Attack on area "Grant"</u>	(battalion unit-13)	<u>Mech unit inside unit-17</u>
(objective-area area19 "grant")	(friendly-unit unit-13)	(represents <polygon:10> (mechanized- infantry-unit unit-10))
(area area19)	(subordinate-maneuver-unit unit-15 unit-17)	<u>Armor unit</u>
(attack-on attack4 area19)	(armor-unit unit-15)	(represents (composite <polygon:6> <ellipse:23> (armor-unit unit-6))
	(battalion unit-15)	
<u>Southernmost friendly task force</u>	(friendly-unit unit-15)	
(composite-unit unit-17)		
(brigade unit-17)		
(friendly-unit unit-17)		

**Figure 5: Sample subset of representations generated for Figure 4**

For the retrieval engine, we used MAC/FAC (Forbus, Gentner, & Law, 1995), a retriever based on the Structure-Mapping Theory of similarity (Gentner, 1983). We chose MAC/FAC because it retrieves cases by mapping systematic sets of similar relations, an approach that favors COADD's (and GeoRep's) structured spatial representations. MAC/FAC can use COADD's descriptions without modification. MAC/FAC also generates candidate inferences from its mapping by noting intersecting but unmapped structure in the old description as potential inferences for the new description. This allows old cases to suggest possible consequences for new cases, since causal links between particular situations and their effects in the old case show up as candidate inferences in the new case when those situations are mapped.

Given a target COA diagram, COADD built a description and MAC/FAC retrieved the most similar COA diagram from a casebase of diagram descriptions previously built by COADD. The description casebase we used contained six division-level COA diagrams with 149-197 visual elements each, and four simpler attack diagrams with approximately 30-50 visual elements each.

Preliminary testing with 10 cases showed that performance was adequate, but not exceptional: similar cases were often retrieved (e.g., for simple attack plans), and the aligned parts were often useful (e.g., the mappings appropriately aligned similar attacks and their associated units). Because the resulting mappings highlighted similar portions of compared diagrams, it was possible to directly show the user which portions of the two diagrams were similar, and delineate those portions as specific correspondences between individual objects in the diagram.

However, as the complexity of the COA diagrams increased, the need to retrieve diagrams based on plan intent and the global unit arrangement became increasingly important. Because COADD could not infer intent, intent could not be used in retrieval. In addition, COADD's simplified domain did not provide enough variability for broader testing. So, while further empirical analysis was

possible, it was decided that any useful test results required a broader subset of the COA symbology.

### The COA Geographic Reasoner

In our second prototype, attempts to expand COADD soon made clear that a deeper revision was needed. A broader symbology and larger diagrams made COADD's recognition difficult and slow. At this time, the COA domain was adopted by the DARPA High-Performance Knowledge Bases (HPKB) initiative as a challenge problem for the upcoming year, and we proposed adapting COADD for this community to create a COA-based "Geographic Reasoner" that would provide geographically-based qualitative spatial relations, along with distance measurements. These relations would provide qualitative geographic relations that would be difficult to construct using standard techniques in an "off the shelf" geographic information system (GIS).

This new role for GeoRep required that it handle a much broader set of COA diagrams than in COADD's simplified domain. Previously, only a dozen or so symbol types needed to be recognized. Now, hundreds of different symbols were possible, a broader set of symbols than that attempted in previous work on symbol-recognition in this domain (Cohen et al., 1997). Handling this larger symbol set required fundamental changes in our approach.

First, we changed the nature of GeoRep's input to use *knowledge-enriched vector graphics*. This format contains primitive visual elements as before, but it also contains information linking visual elements to specific COA objects, in effect pre-classifying those objects. So while previously GeoRep had to recognize armor battalions from primitive elements, it was now told which visual elements were armor battalions. This left GeoRep the more tractable task of representing geographic relations between glyphs.

Using knowledge-enriched vector graphics resolved the scaling issues for symbol recognition. Knowledge-enriched

input also came at minimal cost to the user. Glyph identification was easily incorporated into a COA GUI<sup>1</sup>.

To handle composite elements in the knowledge-enriched input, we added to GeoRep a new *glyph* visual element type. In GeoRep, glyph elements contain a set of component shapes which have display characteristics, extent, and location, but are not analyzed by the LLRD's low-level vision routines.

We also carefully constrained the scope of the Geographic Reasoner's task. In collaboration with HPKB research teams, we determined a set of geographic queries for the Geographic Reasoner (Figure 6). These queries emphasized relative distance measurements, areas of operation, relative direction, and path understanding. Some queries also provided quantitative estimates of some values (e.g., distance along a path between points).

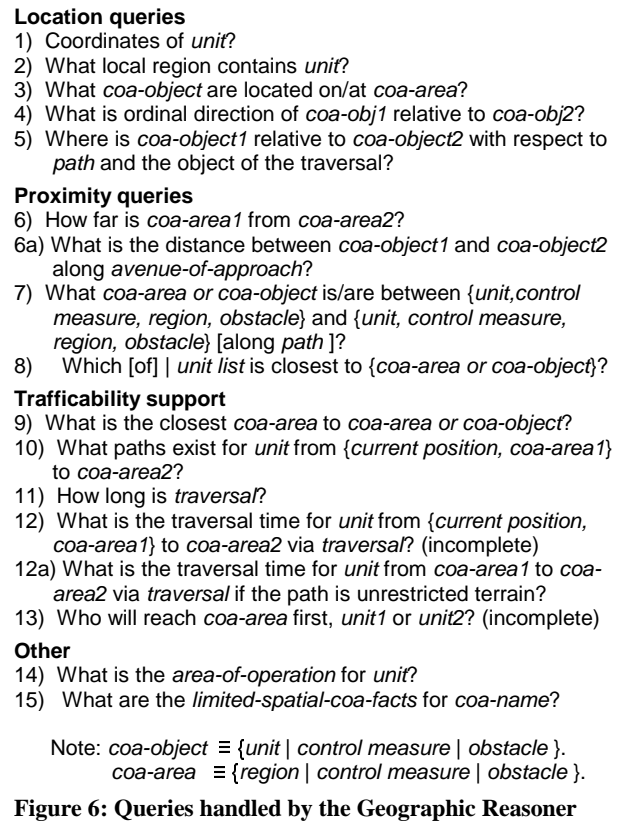
Because the Geographic Reasoner needed to communicate with other reasoners, and also needed access to other reasoners' knowledge, we modified the GeoRep's HLRD to use a different reasoning engine. The earlier system used GeoRep's built-in Logic-based truth maintenance system (LTMS). For this domain, instead of an LTMS, GeoRep built its visual rules using the Domain Theory Environment (DTE). The DTE (Mostek, in preparation) is a reasoning system that has the same functionality as many other theorem-provers, but works from a knowledge base that is saved in a standard ODBC-compliant database.

DTE has the ability to incorporate a much larger set of axiomatic knowledge than the LTMS, and can also exchange this information with other knowledge servers and clients through a KQML socket connection (The DARPA Knowledge Sharing Initiative External Interfaces Working Group, 1993). For geographic reasoning, DTE acted as a host, fielding queries from outside reasoners and passing geometric queries to GeoRep.

**Results.** The Geographic Reasoner answers a broad range of queries, and was successfully used by other research systems as part of their knowledge-based COA critiquers. The ability to use both the imported knowledge about the particular diagram, as well as the ability to access both the conceptual implications of a diagram and its visual content, led to a robust and flexible reasoner.

The reasoner easily handled an extremely large number of queries. In our testing, the reasoner handled 190 geographic queries over four different COA diagrams, and answered all but 8 correctly.

<sup>1</sup> In fact, *two* such COA GUIs were built. One COA diagram builder, constructed by Teknowledge Corporation, used drop-down menus and user dialogs in a standard GIS tool to identify COA symbols as they were placed on a regional map (however, this system could not perform a qualitative analysis of the diagram). Our research group later developed a multimodal sketching system for creating COA diagrams. This system allowed users to request specific object types ("Add armor battalion...") and then note the object location and extent via a pen interface. These systems produced roughly identical representations.



The Geographic Reasoner could also combine different queries to determine a number of critical characteristics in a COA diagram. Figure 7 demonstrates this by showing a small set of questions and answers performed by the Geographic Reasoner for a COA diagram essentially identical to the one shown in Figure 1. In answering these queries, there is a clean interaction between knowledge about the glyph types ("Which glyphs are minefields?"), semantic categories ("Are minefields a kind of obstacle?") and qualitative spatial relations ("Is there an obstacle between this unit and its goal?"). Often spatial relationships turn on conceptual knowledge ("Is this unit a part of the area of operations it is inside? Only if its task does not take it outside of that area.").

While the system is powerful, the resulting visual domain theory is small, containing approximately 51 axiomatic rules and 23 base statements (categories and category relations) covering all the queries in Figure 6.

## Conclusion

COA diagrams utilize a rich form of real-world diagrammatic reasoning that, through its large, composable symbology and its broad use, constitute a useful test bed for research into diagrammatic and spatial reasoning. We have described two systems built in this area using our spatial representation system, GeoRep.

There are several results from this work that have more general applicability. The most important is that qualitative

Here are sample geographic queries, with the Geographic Reasoner's answers in boldface. For brevity, they have been translated into English sentences similar in meaning to the original propositional forms. Units A, B, and C are friendly units located in the dashed region of Figure 1, ordered from top to bottom. Unit Z is the enemy unit to their right.

What is a path between friendly unit B and enemy unit Z?

**Path-880.** Returns a path between the two units along an existing avenue of approach (not shown in Figure 1).

How far apart are the two units along that path? **2.96 km.**

What obstacles are between unit B and enemy unit Z? How far from unit B to the obstacle along that previous path?

What is the ordinal direction from unit B to the obstacle?

**Minefield-84. 2.12 km. East—directly.**

What unit is inside assembly area Tank? **Unit B.**

What is the area of operations for unit A? Answer returns an area, bounded by brigade-level borders, around unit A.

What is the area of operations for unit B?

Answer returns the same area. The returned area is actually located east of the unit. Unit B is not inside this area, but Unit B's task (e.g., represented by the arrow to its right) is inside the area, and the visual domain theory understands that this means that unit B has been assigned to that area.

**Figure 7: Questions and answers to the Geographic Reasoner for Figure 1, re-written in English.**

spatial reasoning techniques can be usefully applied in diagrammatic reasoning. More specifically, a two-level qualitative spatial reasoner, such as GeoRep, which builds an extensible low-level spatial vocabulary from line drawings, provides one mechanism for quickly building reasonably powerful diagrammatic reasoners.

Of course, there are also a number of limitations in both systems described here. The difficulty with scaling COADD for more extensive COA interpretation is telling, and highlights the difficulty of creating spatial reasoners that can make subtle distinctions between a large number of similar glyph types. It remains to be seen whether a more powerful model of low-level spatial reasoning could overcome these difficulties.

While the Geographic Reasoner works well for the set of queries given in Figure 6, this query set is somewhat limited. In addition, performance of the system was often slow, partly due to the slowness of the reasoning engine, but also due to inefficiencies in the LLRD's proximity-detection routines. Queries sometimes took several seconds to complete. Despite these limitations, this system has proven effective in its domain, and is currently being evaluated by the Army for integration into a prototype course of action decision support system.

In the future, we hope to examine possible extensions to the low-level visual vocabulary for COADD in order to clarify what other characteristics lead to scaling difficulties. In addition, we plan to extend the power of the Geographic

Reasoner, increasing the number of query types and speeding up the reasoning in general.

## Acknowledgments

This research was supported by the Defense Advanced Research Projects Agency under the High Performance Knowledge Bases initiative, the Cognitive Science and Computer Science programs of the Office of Naval Research, and by the National Science Foundation under the Learning and Intelligent Systems program. Special thanks to Jim Donlon, who originally proposed this domain. Useful support was provided by Tom Mostek, Jeff Usher, and George Lee. Feedback on the Geographic Reasoner was provided by John Li, Cleo Condoravdi, and David Moriarty. Thanks to LTC John Duquette for information on the history of COA diagrams.

## References

- Cohen, P., Johnston, M., McGee, D., Oviatt, S., Pittman, J., Smith, I., Chen, L., & Clow, J. (1997). QuickSet: Multimodal interaction for distributed applications, *Proceedings of the Fifth Annual International Multimodal Conference* (pp. 31-40). Seattle.
- Department of Defense. (1999). *Department of Defense Interface Standard: Common Warfighting Symbology* (MIL-STD-2525B): United States Army.
- Department of the Army. (1997). *Staff Organizations and Operations* (Field Manual 101-5). Washington, DC.
- Ferguson, R. W., & Forbus, K. D. (2000). *GeoRep: A flexible tool for spatial representation of line drawings*. Paper presented at the National Conference on Artificial Intelligence, Austin, TX.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19(2), 144-206.
- Gentner, D. (1983). Structure-Mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Hendrich, N. (1999). *JavaFIG: The Java diagram editor*, [Web page]. Computer Science Department, University of Hamburg, Germany. Available: <http://tech1.informatik.uni-hamburg.de/applets/javafig/> [1999, December 28].
- Luvaas, J. (Ed.). (1966). *Frederick the Great on the Art of War*. Ontario, Canada: The Free Press.
- Mostek, T. (in preparation) *The Domain Theory Environment*.
- The DARPA Knowledge Sharing Initiative External Interfaces Working Group. (1993). *Specification of the KQML Agent-Communication Language -- plus example agent policies and architectures*.
- Ullman, S. (1984). Visual routines. In S. Pinker (Ed.), *Visual Cognition* (pp. 97-159). Cambridge, MA: MIT Press.