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**
Computer Science Department, Northwestern University
Evanston, IL 60201, USA

[‡]**Battle Command Battle Lab (BCBL)**
415 Sherman
Ft. Leavenworth, Kansas 66027

This paper describes two reasoners built for a real-world diagrammatic reasoning task: Course-of-Action (COA) diagrams. COA diagrams are a useful test bed for diagrammatic reasoning due to their inherently spatial domain and extensible visual symbology. Using a qualitative spatial reasoning engine, GeoRep, we built two COA diagram interpreters. We first describe COA diagrams, then the GeoRep engine, and then the resulting COA interpreters.

COA diagrams are used by the military to depict a region's military units and their assigned tasks (Figure 1). COA diagrams also depict topographical features, movement types (via various arrow types and polyline symbols), available routes, and tasks such as blocking enemy movement. A written plan accompanies the diagram (note: "COA diagram" refers solely to the diagram). COA diagrams are often hand-sketched, and redrawn as needed during planning.

COA diagram interpretation requires two kinds of qualitative spatial reasoning. Locally, symbol shapes indicate object types and characteristics. Globally, relative symbol placement indicates geographically-based relations.

Locally, COAs communicate meaning via their simply-drawn and broadly composable symbology. Symbols are easily classified via their visual structure. Figure 2 shows standard symbols for boundaries, task forces, minefields, friendly and enemy armor battalions, and attacks on objectives. All these symbols use composable subparts. For an armor battalion (Figure 2-d), the rectangle indicates a friendly unit, the contained ellipse indicates armor, and the two "antennae" indicate a battalion. Many parts, such as echelon markers, apply across many symbol classes.

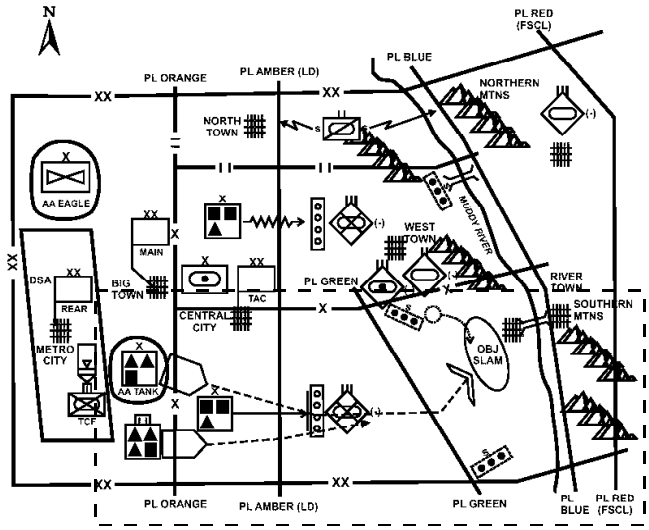

Figure 1: A COA diagram (Dept of the Army, 1997), Figure 5-5). For brevity, we focus on the dashed rectangle area, depicting the main attack (and two supporting attacks) on Objective SLAM.

More globally, COA diagrams also communicate meaning via symbol placement. For example, Figure 1 depicts (in the dashed rectangle) three task forces attacking objective SLAM. Units are assigned to the attack by placement along an attack arrow. Here, the attack paths cross an enemy minefield, and indicate an enemy regiment "behind" that minefield. The diagram boundaries divide the map into "areas of operation" to which units are assigned.

We built our COA interpreters using a qualitative spatial representation engine, GeoRep (Ferguson & Forbus, 2000). As input, GeoRep takes a drawing in a vector graphics format. From this, it creates a qualitative spatial representation of the drawing in a domain-specific spatial vocabulary.

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 drawing, representing primitive visual relations such as proximity, parallel line segments, polygons, and connection and containment relations. These relations are salient in early vision. The HLRD in turn uses domain-specific rules to extend the LLRD's representation to detect domain-specific spatial relations and symbols types, which are given as output.

The COA diagram describer

Our first reasoner is the *COA diagram describer* (COADD). COADD uses GeoRep to describe COA diagrams containing a simplified symbology. Its symbol set includes assembly, engagement, and objective areas, basic unit and attack types, and borders.

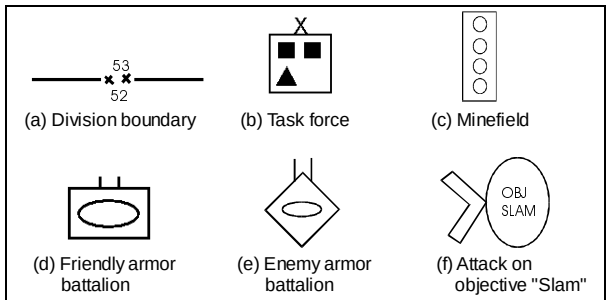

Figure 2: Typical COA symbols used in Figure 1

HLRD rules are used to recognize COA symbols using low-level relations. For example, friendly armor units are recognized as rectangles containing horizontally-oriented ellipses. These rules reflect the compositionality of COA symbols. For example, echelon markers are detected and then linked to specific units and boundaries, reflecting how echelon markers apply across different symbol types. Other rules infer unit intent by proximity to attack arrows and assembly areas. The rule set is small but expressive, with 37 HLRD rules covering 18 object types and relations.

To test COADD, we built a COA retriever. We used MAC/FAC (Forbus, Gentner, & Law, 1995), a retriever based on the Structure-Mapping Theory of similarity (Gentner, 1983), as the retrieval engine. Given a target diagram, COADD built a description and MAC/FAC retrieved the most similar COA diagram, using a casebase of previously-built descriptions.

Preliminary testing with a casebase of 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, but the depth of COADD's simplified domain did not provide enough variability for proper testing. Useful test results required a broader subset of the COA symbology.

The COA Geographic Reasoner

In our second prototype (which was designed for the DARPA HPKB initiative), attempts to expand the COADD prototype soon made clear that a deeper revision was needed. A broader symbology and larger diagrams made COADD's recognition difficult and slow.

First, GeoRep's input was changed to use *knowledge-enriched vector graphics*. This format contains primitive visual elements as before, but also links visual elements to specific COA objects—in effect pre-classifying them. Identified symbols are handled by a *glyph* visual element, which contains component shapes that have display characteristics, extent, and location, but are not analyzed by GeoRep's low-level vision routines. For example, while GeoRep previously had to recognize armor battalions, the input now specifies which visual elements are armor battalions. This leaves the more tractable task of representing geographic relations between glyphs.

In collaboration with other HPKB research teams, we determined a set of 15 geographic queries. These queries emphasize relative distance and direction, areas of operation, paths, and metric distance measures.

The resulting geographic reasoner answers a broad set of queries, and was used successfully by other research teams to build a knowledge-based COA critiquer. In our testing, the reasoner handled 190 geographic queries over four different COA diagrams and answered all but 8 correctly. While the system is powerful, the visual domain theory is small, containing 51 axiomatic rules and 23 base statements (categories and category relations) for 15 query types.

The geographic reasoner combines spatial and semantic knowledge to determine critical relationships. Figure 3 shows several of its queries and answers. Its answers show a clean interaction between knowledge about the glyphs ("Which glyphs are minefields?"), semantic categories ("Are minefields an 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's goal is not elsewhere.").

Units A, B, and C are friendly units in the dotted region of Figure 1. Unit Z is the enemy unit to their right.

What is a path between friendly unit B and enemy unit Z? How far apart are they along that path? **Path-880** *(a path between the units).* **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 kilometers**. **East—directly**.

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? *Same area, located east of the unit. Unit B is outside this area, but Unit B's task is inside, meaning that it is unit B's responsibility.*

Figure 3: Questions and answers to the Geographic Reasoner for Figure 1, re-written in English.

Conclusion

COA diagrams constitute a useful test bed for research into diagrammatic reasoning. Using a two-level qualitative spatial reasoner like GeoRep, one can quickly build powerful diagrammatic reasoners.

There are many limitations in the current prototypes. The difficulty with scaling COADD is telling, and highlights the difficult task of distinguishing between many similar glyph types. While the Geographic Reasoner works well, its query set is limited. In addition, performance is sometimes slow due to inefficiencies in proximity-detection.

Development continues. We hope to extend the low-level visual vocabulary to clarify COADD's scaling difficulties. We also plan to extend the Geographic Reasoner, adding query types and increasing reasoner efficiency. This system is currently being evaluated for integration into a prototype COA decision support system.

Acknowledgments. This research was supported by the DARPA High Performance Knowledge Bases initiative, by the Office of Naval Research's Cognitive and Computer Science programs, and by the National Science Foundation under the Learning and Intelligent Systems program.

References

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.* AAAI-2000. 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.