

AFRL-IF-RS-TR-2004-339 Vol 2 (of 2)
Final Technical Report
December 2004



FACILITATING SUBJECT MATTER EXPERT (SME)-BUILT KNOWLEDGE BASES (KBS)

Information Extraction & Transport, Incorporated

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K180

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-339 Vol 2 (of 2) has been reviewed and is approved for publication.

APPROVED: /s/

CRAIG S. ANKEN
Project Engineer

FOR THE DIRECTOR: /s/

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE DECEMBER 2004	3. REPORT TYPE AND DATES COVERED Final Jul 00 – Apr 04	
4. TITLE AND SUBTITLE FACILITATING SUBJECT MATTER EXPERT (SME)-BUILT KNOWLEDGE BASES (KBS)			5. FUNDING NUMBERS C - F30602-00-C-0173 PE - 62301E PR - RKFM TA - 00 WU - 02	
6. AUTHOR(S) Edward J. Wright				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Information Extraction & Transport, Incorporated 1911 North Fort Myer Drive, Suite 600 Arlington Virginia 22209			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFED 3701 North Fairfax Drive 525 Brooks Road Arlington Virginia 22203-1714 Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-339 Vol 2 (of 2)	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Craig S. Anken/IFED/(315) 330-2074/ Craig.Anken@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This report addresses the functional performance task of classifying terrain according to its degree of (un)suitability for the performance of a given activity (e.g., locomoting, shooting) by a given military unit/vehicle, under given tactical conditions (e.g., proximity/location/type of enemy forces, susceptibility to enemy intelligence collection assets, weather). Such a terrain modeling/reasoning capability serves as foundational infrastructure for battlespace reasoning applications including information fusion, planning, and simulation applications. The goal of this work was to exploit the Knowledge Representation and Reasoning (KR&R) tools in a synergistic architecture with logical reasoning capabilities, and develop a proof of concept prototype terrain suitability knowledge base.				
14. SUBJECT TERMS Knowledge Base, Knowledge Representation, Capture, Challenge Problem Development, Terrain Suitability				15. NUMBER OF PAGES 60
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1.	Introduction.....	1
1.1	Organization of the document.....	2
2.	Background.....	3
2.1	Detailed Terrain Data.....	3
2.1.1	Standard Terrain Analysis Applications.....	3
2.1.2	Military GIS Data	5
2.1.3	Environmental Data Coding System (EDCS).....	9
2.1.4	Complications in the use of geospatial data.....	11
2.2	Current Terrain Analysis Practice.....	12
3.	Probabilistic Terrain Assessments	14
3.1	Representing Terrain Analysis Models as Probabilistic Models	16
3.1.1	Quiddity*Suite	19
4.	Terrain Suitability Knowledge Base.....	22
4.1	Logical Reasoning	23
4.2	Logic Models	24
4.2.1	Unit Models	24
4.2.2	Activity Models	24
4.3	Probabilistic Models	25
4.3.1	Data Quality Models.....	25
4.3.2	Missing Data Models Vegetation Models.....	27
4.3.3	Terrain Assessment Models.....	31
4.4	Knowledge-Based Model Construction Implementation.....	33
4.5	Scenarios.....	40
4.6	Application.....	42
5.	Extension to Support FFD or Other Terrain Data Sets.....	43
5.1	Expanded Range of Terrain Data Formats.....	43
5.2	Expanded Range of Terrain Data Content.....	43
5.2.1	Interoperable Definition of Terrain Features and Attributes	45
5.3	Tasks to Extend TSKB to Support FFD.	45
6.	Conclusions.....	46
7.	References.....	47

Appendix A. Acronyms & Expansions.....	50
Appendix B. User Recognition of the Need for Quality / Uncertainty Assessments	52

List of Figures

Figure 1. Line of Sight (LOS) Product. Left - shaded relief view of an experimental high resolution elevation data set (1 meter resolution), with the location of an observer (blue triangle). Right - traditional LOS product display, which identifies areas that can be observed (green) and areas that are obscured by terrain(red).....	4
Figure 2. Cross Country Mobility (CCM) product. This display shows the predicted CCM speed of an M1 tank for a small area of Korea, based on the DMA mobility model and ITD data.....	5
Figure 3. Different views of a Digital Elevation Model.....	6
Figure 4. Information content of Interim Terrain Data (ITD).....	8
Figure 5. Traditional CCM display generated without concern for quality of terrain data.....	14
Figure 6. A probabilistic CCM display that includes the uncertainty in the terrain data.....	15
Figure 7. Prototype Bayesian Network for propagation uncertainty through a GIS Model.....	17
Figure 8. Bayesian Network the encodes the ETL CCM algorithm.....	18
Figure 9. Decision Model for the standard application of terrain analysis in COA analysis.....	18
Figure 10. Decision model for application of terrain analysis when the commander understands the uncertainty of the terrain analysis results.....	19
Figure 11. The components of the Terrain Suitability Knowledge Base.....	22
Figure 12 Overview of First Order Logic system use.....	23
Figure 13. Basic data quality models.....	25
Figure 14. A more complex error model that includes relationships between several feature types.....	26
Figure 15. An error model constructed from several Bayesian Network fragments.....	26
Figure 16 Natural Vegetation Regions of the World.....	27
Figure 17. Example of inferred vegetation data for terrain reasoning.....	28
Figure 18. The Earth can be divided into regions called biomes, with similar geology, climate and vegetation. Geographic relationships are expected to be similar with biomes.....	29
Figure 19. A complex relationship model.....	30
Figure 20. An additional terrain feature extends the relationship model.....	31
Figure 21. Example probabilistic CCM product.....	31
Figure 22. Line of Sight Model.....	32
Figure 23. Probability of LOS.....	33
Figure 24. Example Probabilistic LOS products, generated with alternate relative error models.....	34
Figure 25 Context and Vehicle Frames.....	35
Figure 26 True Terrain and Cross Country Mobility Frames.....	36
Figure 27 Example of parent-child slot relationships.....	36
Figure 28: Graphical representation of links between frame slots.....	37
Figure 29. CCM frame example.....	37
Figure 30. Representing relations between frames requisite for causal influence.....	38

Figure 31. Scenario Example: Ambush with intent to capture supplies.....	40
Figure 32. Example of terrain suitability factors.....	41
Figure 33. The application of the TSKB as part of an automated COA evaluation capability.....	42
Figure 34. A Bayesian Network for the CCM algorithm, showing the portion that changes based on the data, and the portion that is independent of the data.....	44

1. Introduction

This report describes accomplishments made by Information Extraction & Transport (IET¹), Inc. on U.S. Government contract F30602-00-C-0173 during the period July 10, 2003 – May 9, 2004. IET's Y3 KB addresses the functional performance task of classifying terrain according to its degree of (un)suitability for the performance of a given activity (*e.g.*, locomoting, shooting) by a given military unit/vehicle, under given tactical conditions (*e.g.*, proximity/location/type of enemy forces, susceptibility to enemy intelligence collection assets, weather). Such a terrain modeling/reasoning capability serves as foundational infrastructure for battlespace reasoning applications including information fusion, planning, and simulation applications.

Cyc and its associated KRAKEN toolset, like other RKF technologies, have not so far been developed with information fusion applications or entailed probabilistic reasoning in mind. IET has (so far independently of RKF) developed extensive, fusion-oriented KR&R tools² integrating probability with the (restricted first-order) logic of frames. The goal of the Y3 effort is broadly to exploit the KR&R tools in a synergistic architecture with logical reasoning capabilities, and develop a proof of concept prototype terrain suitability knowledge base.

In negotiations during the summer of 2003, three specific IET tasks were specified:

1) Produce a hybrid (logical and probabilistic) Tactical Terrain Reasoning engine with:

- two terrain factors: Cross Country Mobility (CCM) and Line Of Sight (LOS), with a capability to interact with commercial GIS software and access basic National Geospatial-intelligence Agency (NGA) terrain Data (Digital Terrain Elevation Data - DTED, Interim Terrain Data ITD & Vector Product Format (VPF) ITD - VITD)

- a Knowledge Base for relevance reasoning for the two initial terrain suitability factors, CCM and LOS. Relevance reasoning is to determine, based on the current situation, when these factors are important. Further, based on the current context, it is to determine the relevant terrain features and attributes.

2) Develop Knowledge Based inference of critical tactical terrain attributes required for COA tactics estimation. This inference supports:

- meta reasoning (about part-whole and isa relationships) and application of geographic analogs to identify substitute inferences when the specific terrain features and attributes required are not available,

- data learning: interaction with external databases to fill in parameters of probabilistic models

3) Develop a design for integrating NGA's FFD, the standard NIMA terrain data product most likely to be available in future conflicts. Because FFD is not completely populated with features and attributes required by specific terrain suitability models, this includes extending the relevance reasoning, meta reasoning, and data learning to handle this product.

¹ A table of acronyms and expansions appears in Annex A.

² Also known as (AKA) Quiddity*Suite

1.1 Organization of the document

Section 2 provides background information on terrain data products used in military planning and decision-making, the current processes by which these products are provided and the terrain data sets that are used to produce them.

Section 3 provides an overview of probabilistic assessment of terrain.

Section 4 details IET's Terrain Suitability Knowledge Base (TSKB), describes the technology that provides it functionality, and presents some examples.

Section 5 identifies the required developments to extend the TSKB to be able to exploit FFD or any other terrain data source.

Section 6 provides conclusions.

Section 7 contains references.

Annex A is the list of acronyms used in this document.

Annex B provides additional information that supports the importance of terrain suitability assessments to intelligence data fusion and IPB, as well as evidence of growing understanding that including probabilistic assessment of uncertainty is important to terrain suitability assessments.

2. Background

During RKF year 3, IET activities were focused on technology development to meet specific needs of the RKF program. Technical needs include integration of logical and probabilistic reasoning, and development of terrain knowledge bases to support automated reasoning about military courses of action. The latter requires a detailed understanding of terrain, including effects of uncertainty in terrain data quality, in order to understand how terrain will constrain military operations, as well as how terrain may offer opportunities, and present risks, to military operations. As a result, IET's technology development is focused on development of a proof-of-concept Terrain Suitability Knowledge Base (TSKB), which has application to reasoning about military courses of action. The TSKB will access detailed terrain information accessed through a commercial Geographic Information System (GIS), a probabilistic representation of terrain data quality and predicted terrain effects, and will also integrate logical reasoning to exploit existing knowledge in logical databases (such as the Cyc KB).

2.1 Detailed Terrain Data

Reasoning about military courses of action must include a critical assessment of the effects of the environment (terrain and weather) on the military operations proposed. The current RKF capability, which uses terrain sketches produced by Nu-Sketch, may be suitable for capturing high-level strategic concepts; however it does not provide the detail needed for operationally realistic terrain reasoning. The US Army and other services have recognized the need for detailed terrain analysis to support military planning and operations by their requirements for detailed terrain analysis data and their emphasis on terrain analysis support for all military operations. State of the Art military reasoning is still done by humans, but is supported by manual and automated terrain analysis assessments produced from detailed terrain analysis data products. Terrain data is manipulated by fielded terrain analysis systems, like the Digital Topographic Support System (DTSS), that use commercial GIS software.


IET's TSKB is was built to exploit modern terrain data products through interaction with the same commercial GIS software used in the DTSS system. The initial TSKB will exploit Digital Terrain Elevation Data (DTED), Interim Terrain Data (ITD), Vector Product Format (VPF) Interim Terrain Data (VITD), and can be enhanced to support Foundation Feature Data (FFD) and other terrain products in the future.

2.1.1 Standard Terrain Analysis Applications

Geographic Information Systems (GIS) have received broad acceptance in a wide range of military applications, supporting decision making during military planning and operational command and control (military applications of GIS are often called Terrain Analysis or Terrain Evaluation). The utility of these applications has created a large demand for geospatial data to support them. Unfortunately, the demand for geospatial data has exceeded the ability of production agencies to produce data; as a result geospatial data from a wide variety of sources is being used, often with little regard to the data quality. A concern is the influence of errors or uncertainty in geospatial data on the quality of military decisions made based on displays of geospatial data.


Two common military applications of GIS are Line of Sight (LOS) products, and mobility products. There are a large numbers of additional GIS products used in military applications, but these two provide a representative sample that is suitable to illustrate challenges of assessing suitability.

The LOS and mobility products are examples of military Tactical Decision Aids (TDA) that predict the effects of terrain on military operations. They are intended to provide relevant information to military decision makers, without requiring them to be experts in geospatial data or the techniques of geospatial data analysis.



A LOS product is based on elevation data, and is used to determine if a line of sight exists between specific points in space. The product may be a terrain profile between two points, or it may be a two dimensional display showing areas that are visible from a defined point and areas that are blocked by terrain. An example is shown in Figure 1. In this example, the observer is on the ground and the product shows areas of the ground that can be observed. Other LOS products might be based on aerial observers at some defined altitude. The LOS product is used by military decision makers to place surveillance systems (observation posts or radar systems), predict the coverage of airborne sensors and to locate direct fire weapon systems. The traditional LOS display shows an absolute, deterministic prediction - without any estimate or visualization of the influence of the potential errors of the terrain elevation data on the result.

An example of a mobility product is the Cross Country Mobility (CCM) product, which is a graphic display of the capability of the terrain to support the off road movement of units equipped with a specific type of vehicle. An example is shown in Figure 2. CCM products are produced from feature data which contains information about terrain soil types, surface roughness, vegetation, and slope (which may be derived from elevation data). There are several CCM algorithms in use - the CCM product in Figure 2 was produced using the DMA CCM algorithm (DMS, 1993). CCM products can be generated for specific vehicle types, for classes of vehicles, or for military unit types. The products can be generalized to produce mobility corridors, or combined with other information to generate avenues of approach for friendly or enemy forces. The traditional CCM display, which may be a hardcopy product or a computer



graphic, shows predicted speeds without any attempt to estimate or communicate the quality of the prediction based on the quality of the underlying data and the quality of the algorithm (GIS model) used to make the prediction.

2.1.2 Military GIS Data

There are a wide range of military digital mapping products (digital terrain data) available from the DoD National Geospatial-Intelligence Agency (NGA). Two classes of data products that represent those most commonly used in military GIS analysis - Terrain Analysis, are the Digital Elevation Model (DEM), and feature data.


There are a number of different ways to encode an elevation surface in a digital file and in general, any of them may be called a DEM. The most common way is as an array of elevation values. Elevation values are provided on a grid with a defined spacing in the North-South and East-West directions. The grid spacing is a measure of the resolution of the DEM: smaller grid spacing corresponds to a higher resolution.

Figure 3 shows a representation of an elevation surface as a grid of elevation values, a contour map and as a three dimensional surface.

A standard DEM product produced by NGA is the Digital Terrain Elevation Data (DTED). NGA produces DTED level 1 data in cells covering an area of 1 degree by 1 degree, with a grid spacing of 3 arc seconds (approximately 100 meters at the equator). DTED level 2 is produced over smaller areas with a grid spacing of 1 arc second (approximately 30 meters at the equator) (NIMA 1996). Specifications for higher resolution DTED at levels 3, 4, and 5 are under development. DTED is widely used for visualization and LOS applications.

Feature data provides information about characteristics of the Earth or objects on the Earth. A wide variety of feature data products have been produced, and are in use for military terrain analysis applications. Most have similar feature content to Interim Terrain Data (ITD).

ITD is a widely available digital feature data in use by military GIS systems today. It was originally developed as an interim product, while users awaited a more detailed and robust digital terrain data product. ITD is available in two forms - ITD, and VITD (Vector Product



Format (VPF) ITD) - that differ in format, although most of the information content is similar. ITD is digital vector data, where terrain features are represented as points, lines and polygons. Each terrain feature has a number of feature attributes defined for it. Figure 4 shows a graphic that illustrates the information content of ITD. Information is provided in 6 thematic layers.

Each layer contains features as points, lines, or polygons, and has an associated set of feature tables that contain attributes of each feature. Vegetation polygons are defined for several types of wooded areas, orchards, and agricultural applications. Vegetation attributes include vegetation stem spacing, and stem diameter. The transportation layer contains features that represent roads, bridges, railroads, airfields, etc. Attributes define road widths, construction materials, bridge length, width, capacity, etc. The surface materials layer provides polygons of soil type and an attribute for surface roughness. The surface drainage layer contains information on rivers and streams, with attributes that define width, depth, bank height and slope. Surface configuration layer contains polygons for surface slope in defined categories. The obstacle layer contains information of other terrain features (like ledges, fences, pipelines, cuts and fills) that may be obstacles to military mobility (NIMA 1996).

ITD is used for a range of military GIS applications (Terrain Analysis) including mobility products like CCM. ITD data, and other feature data products are very valuable, they are also expensive to produce, requiring lots of human intensive feature extraction. NGA has recognized the inability to provide widespread coverage of ITD (or ITD like data) in support of worldwide military operations. The NGA concept for future terrain data support envisions large area coverage of a subset of quickly produced data (Foundation Feature Data - FFD) to meet the military's immediate planning needs, and a capability for rapid production of more complete data (Mission Specific Data SETS - MSDS) to meet specific requirements identified once a crisis starts.

US military's priority areas of interest now extend over the entire Earth. Because of the high cost of producing feature data, high or even moderate resolution feature data are not universally available. The new NGA production processes that are generating FFD and MSDS datasets have not caught up with the demand. The result is that for any operational area there is, in general, no uniform spatial data coverage. Most areas are covered by low resolution, wide area data. For some areas, there is medium quality data available, and for limited areas there may be patches of even higher quality data. During a military operation, the available geospatial data will grow rapidly as NGA and other production centers generate data in response to military requirements. But at any time – the geospatial database will be a heterogeneous mixture of data types, resolutions, with different currency and accuracy.

Quality of geospatial data is an issue that has received considerable interest in the academic GIS community (Goodchild 1992). Studies have shown that, while all geospatial data contain errors, errors in geospatial data are not well documented, not well understood, and are commonly underestimated by users. Military geospatial data organizations have shown considerable interest in establishing specifications for data, and in evaluating data sets to ensure that they meet the prescribed standards. However, until recently military GIS operators and users have shown little interest in understanding and managing uncertainty in geospatial data for military applications.

develop experience using only the standard data. In the real world – with constrained production resources, and where crisis arise quickly in unpredictable locations – military decision makers are faced with planning and conducting operations before the standard data sets can be produced. Faced with operational imperatives, military decision-makers will use whatever geospatial data are available. While it is good that US military systems are (usually) flexible enough to exploit available data, there are significant risks. Because all automated command and control systems have been designed to use the standard data, these systems will likely have built in assumptions about the quality of the data. These assumptions are unlikely to be explicitly documented. When these systems are employed using non standard, available data, in a crisis –the user will likely assume that the displayed results are just as good as the standard data previously experienced. As a result, the user may make inappropriate decisions.

This problem does not go away if we assume that users will always insist that the standard data be produced and provided to them before they take action. Recent history is full of examples where military actions were required, and taken, long before the standard data sets were available. There is a challenge in the other direction as well. During sustained military operations, considerable data generation resources are dedicated to producing geospatial data to support operations. In time all the standard data sets are available. In addition, special high-resolution data sets have often been made available. If systems and decision makers are only accustomed to using standard data, there is no way to appropriately exploit the higher resolution and higher accuracy of the special data sets.

A potential solution to this challenge is to develop standards for data quality metadata for all military geospatial data. Automated systems could be built that are capable of reading the data quality metadata, propagating the data uncertainty through the various TDA models into a prediction of the uncertainty in the TDA product, and displaying the uncertainty in some usable way to the decision maker. Then the automated systems and their users could use any dataset that meets the standard for data quality metadata with confidence.

2.1.3 Environmental Data Coding System (EDCS)

Terrain Data Content

Terrain data is produced in a wide range of terrain data products by a large number of government and civilian organizations in many countries. The profusion of terrain data products can result in confusion over the meaning of the content of the data. The content challenge is based on the meaning of the labels used to describe features or attributes in the data set. Data produced by different organizations, for different purposes, may use the same label as a feature class but have very different meanings. Even a term as simple as “road” could mean very different things in different data products. Because of this issue, many of the terrain analysis software capabilities that are available will only work with a specific terrain data product.

Within the military Modeling and Simulation (M&S) community, the SEDRIS organization (www.sedris.org) was formed to address this issue, and has developed the Environmental Data Coding System (EDCS), to address the challenge of interoperability of terrain data content.

From the SEDRIS web page:

Environmental data is an integral part of many of today's information technology applications. The use of environmental data will grow substantially as availability and

access to such data increases, and as tools for manipulation of environmental data become less expensive and more sophisticated.

As this trend continues, the representation and sharing of environmental data will play a key role in the interoperation of heterogeneous systems and applications that use such data. This need was recognized in the mid-1980's, when the ability to network large numbers of heterogeneous simulation systems became a practical reality. Research and work in this area continued while a better and more complete understanding of the complex issues associated with describing and sharing of environmental data for a wide variety of (simulation) applications was formed. SEDRIS was conceived in order to tackle these issues in a uniform and unified manner.

Although the initial application domain for SEDRIS stems from the needs of the modeling and simulation field, it was immediately recognized that the representational technologies required to capture and communicate environmental data are fundamentally one and the same and, in large part, can be dealt with independent of the end-applications.

At the same time, it was also understood that too often end-applications shape and form the characteristics of how data and data representation are used. The challenge for SEDRIS was to provide a means for representation and sharing of environmental data that not only was efficient in practical use, but also was specific enough to address the real needs of a wide variety of end-applications, and at the same time preserve the degree of semantics needed for others to understand the nature of the data. The range of end-applications included representation of environmental data for such applications as analysis, visualization, simulation, planning, modeling, etc. This took into account the meteorological and oceanographic communities, the simulation sector (both military and commercial), the GIS (or more broadly, the environmental information systems) community, the military operational community (i.e., C4I), as well as others who needed to share or communicate environmental data.

Added to this was the goal of getting away from stovepipe views of the environment, and providing a mechanism that also allowed for integrated environmental data to be represented. Integrated environmental data, where ocean, terrain, atmosphere and space data (about a region) can be seamlessly represented, was recognized as a key component of many future information technology applications. And although very few applications today deal with such diverse data at the same time, developers of SEDRIS believed such a need would be a reality in the future. (http://www.sedris.org/ab_1trpl.htm)

The EDCS provides a mechanism to specify the environmental "things" that a particular data model construct is intended to represent. That is, a "tree" could be represented alternatively as a <Point Feature>, an <Aggregate Geometry>, a <Data Table>, a <Model>, or some combination of these and other data modeling constructs. Which of these the data modeler (i.e., the data provider of a SEDRIS transmittal) chooses is orthogonal to the semantic of the "thing" that is represented (and its location). The provision of such a "thing" in a SEDRIS transmittal pre-simulation must result in a shared understanding of "what the thing is and what it potentially means" to all participating applications.

In addition, the EDCS provides mappings between alternate representations of terrain "concepts" (features and attributes) used in standard terrain data products and the EDCS. For example, there

is a mapping between the Feature and Attribute Coding Catalogue (FACC), used in many NGA products, and the EDCS.

2.1.4 Complications in the use of geospatial data.

There are a number of technical factors that complicate the use of geospatial data for terrain analysis applications. Complications result from diverse formats of elevation and feature data, differences in data content, and from differences in research.

2.1.4.1 Data Formats

Digital geospatial data are available in a wide range of digital formats, and must usually be converted into a common format before the data can be used in analysis. The first format challenge is physical media format, the second in logical format. There are examples of organizations being unable to exchange geospatial data, even after agreeing on the specification, digital format, and hardware medium, because they were using different versions of the computer operating system and the default block size for the tape drive changed from between versions.

The next potential challenge is logical format. There are a wide range of government and commercial data formats for both raster and vector geospatial data. NIMA produces data in standard formats for military command and control system. These formats are supported by most commercial GIS software. The logical format challenge usually arises when non-standard data sets are being exploited.

2.1.4.2 Coordinate Systems

All spatial data by nature include information about location. Location is specified by spatial coordinates in a defined coordinate system. The horizontal coordinate system requires specification of the geographic or geodetic datum, and the mapping projection used. The datum specifies the origin and orientation of the coordinate system, and the size and shape of the reference ellipsoid used as the Earth model. Historically, datums were defined locally using astronomic observations. As a result, the location of a point defined by different datums may differ by hundreds of meters. The vertical datum defines the zero value of the elevation scale. Mean sea level is widely used in mapping, although differences in historical surveys that defined mean sea level can result in differences between different “mean sea level” datums. An alternative vertical datum, is ellipsoid height: the height above the ellipsoid model of the Earth. Use of different ellipsoids, and origins of the ellipsoid can result in large differences in ellipsoid heights. NIMA has identified over 100 geodetic datums in common use around the world today. Most new NIMA products are being produced using the World Geodetic System (WGS) datum. Standard software is available to convert between datums. Failure to recognize that locations in geospatial data are recorded in different datums, and to convert all data to the same datum, can result in serious errors in analysis.

In addition to different datums, geospatial locations are specified using different map projections. Map projections provide a mathematical transformation between the curved surface of the Earth (or the ellipsoid model of the Earth) and a flat piece of paper. Even today with digital processing by computers, the “flat” representation provided by the map projection provides computational advantages and is widely used. For example, with most map projections, computations can be

performed using plane geometry instead on spherical or ellipsoidal geometry. Converting all data sets to the same projection is an important processing step for geospatial analysis. All common GIS software systems provide capabilities to convert between different map projections.

It is always possible to convert data from one coordinate system to another, by performing a datum conversion and/or reprojecting onto a different map projection. This process may introduce errors. For example, any conversion from one datum to another, will involve using parameters (derived by a variety of means) that are subject to error. As a result, the new coordinates, in the new datum, will all have additional error as a result of the conversion. This error will show up as an unknown bias for all points close to each other. NIMA provides estimates of the accuracy of datum conversion parameters and a means of estimating the accuracy of the converted coordinates.

Unlike a datum conversion, the conversion of data to another map projection uses an “exact”³ mathematical formula. However, this process may still introduce error in some types of data. Vector data, which has coordinates of specific points, can be re-projected without error. For raster or gridded data, the re-projection process requires estimating values at new raster or grid points. A number of methods are commonly used: nearest neighbors, bi-linear interpolation, convolution, or many alternatives. The best choice of resampling method depends on the intended use (Schowengerdt 1997). In every case, the potential for introducing error should be considered and documented in the metadata for the resampled product.

2.2 Current Terrain Analysis Practice


Terrain analysis products are produced as part of the Intelligence Preparation of the Battlefield (IPB) process. In the Army and Marine Corps, there are dedicated topographic / terrain analysis units equipped and trained to provide terrain analysis support to commanders and their staffs. These units are equipped with the Digital Topographic Support System (DTSS) with computer hardware and software, including Geographic Information Systems (GIS) and image processing software, that provides capabilities to manipulate standard and nonstandard terrain data to generate terrain analysis products to support military operations. The topographic soldiers who provide the terrain analysis support are using automated tools, but their training and experience provides appreciation for the limitations of the terrain data and of the algorithms they are using. They are also trained to interpret the results in the context of military operations that they are supporting. However their capabilities do not include the ability to quantitatively assess the impact of terrain data quality on the accuracy of the products they produce.

There are also a number of software packages that provide basic manipulation of terrain data, which are now available on standard military command and control computers. Examples include ArcView, Terrabase, and Falcon view. These software packages have made it possible for almost anyone in the military with access to a computer to produce their own custom terrain products. There is a draw back to this wide availability of terrain analysis software, because some of the users are not aware of the limitations of the data and of the algorithms they are using, and so may draw inappropriate conclusions from the products that they produce.

³ Some map projection equations are approximations derived by series expansions, which are truncated at some number of terms. It is always possible to include additional terms to achieve any desired accuracy (Snyder 1987).

In addition, the current trend is to automate more and more of the planning functions in military mission planning and command and control systems. These automated functions include algorithms to produce the standard terrain analysis products (like LOS and CCM) to support COA generation. This trend began with the development of automated planning capabilities to support M&S. These capabilities are being used in operational Mission Planning and Rehearsal Systems (MPRS), and are being transitioned to operational command and control systems. The potential drawback is that automated systems will be generating planning and operational options for commanders, when the commander does not understand the limitations of the data and of the algorithms they are using, and so may draw inappropriate conclusions from the presented options.

3. Probabilistic Terrain Assessments



Unfortunately, exploiting detailed terrain data is not enough for a thorough terrain analysis assessment. Even the highest resolution, and most current data is not 100% accurate and many of the terrain analysis models are extremely sensitive to even small variations in terrain data. Failure to take into account the influence that uncertainty in the terrain data can have on predictions of the impact of terrain on military operations can result in missed opportunities or dangerous surprises. Figures 5 and 6 present an example. In Figure 5, a traditional Cross Country Mobility (CCM) product is shown. CCM predicts the speed that a specific vehicle (or unit composed of a specific vehicle) can operate across country (off roads). The CCM product uses slope, vegetation, ground roughness, soil and soil moisture data. The legend shows predicted speed ranges that are color coded on the CCM product: red for No Go terrain, green for Go terrain, and several intermediate colors for Slow Go terrain. CCM is a common terrain analysis product used to support analysis of avenues of approach and courses of action. Figure 5 also shows a defensive

obstacle belt (blue triangles) designed to tie into natural terrain obstacles (No Go terrain) to block the potential approach of an enemy unit. This is a doctrinally correct plan, based on the information available. Figure 9 shows the same area with a CCM product that does include the uncertainty in the terrain data. The legend again maps the predicted speed range to the same colors, however in this product a bi-variate legend is used. The quality of the color represents the uncertainty in the prediction. Bright colors are areas where the prediction is believed to be fairly accurate, while the darker colors are areas where there is considerable uncertainty in the CCM prediction. Also shown in Figure 9 is a “drill down” into an area of the map. This area is dark red, indicating that while the most likely CCM is “No Go” there is considerable uncertainty in this prediction. The “drill down” shows specific probability distribution across speed ranges. While the highest probability (28%) is No Go, there is also a fairly large probability of fairly fast go (Probability that CCM > 20 Kph is 52 %).

This bi-modal distribution of predicted CCM speeds can happen fairly easily. It is caused by uncertainty in the soil type, combined with uncertainty in the soil moisture. Some soil types (silt and clay) break down quickly when they are wet – and will not support vehicle movement. Other soil types (sands and gravels) do not break down when wet, and can continue to support vehicle traffic.

The bi-modal distribution reveals that there is significant risk that the area believed to be No Go, may be a Go area. If that turns out to be true there is an opportunity for the opposing force (Red) to use this area as an avenue of approach which bypasses the Blue obstacles. A Blue


commander, who is aware of this risk can take actions to mitigate the risk: prepare additional obstacles, task sensors or reconnaissance to collect additional terrain information, or employ sensors or observations posts to detect enemy activities in this avenue in time to react to them. These mitigating actions are not possible with the information in Figure 9.

Figures 5 and 6 serve as one example of the pitfalls of ignoring uncertainty in terrain data and the resulting uncertainty in assessments of the impact of terrain on military operations. IET's TSKB is being built using probabilistic representations of terrain data quality and probabilistic assessments of the impact of terrain based on the quality of terrain data. The probabilistic result reveals potential risks and opportunities that are not revealed by traditional terrain analysis methods.

New sensors for High resolution data collection

Some will argue that the current and future availability of new sensors like Interferometric Synthetic Aperture Radar (IFSAR) and Light Detection and Ranging (LIDAR) sensors capable of rapidly generating high resolution and high accuracy terrain data make the issues of data quality and uncertainty less important. There are several problems with this view. First the availability of high resolution and/or high accuracy data will remain limited. These sensors are commonly mounted on fixed wing aircraft or Unmanned Aerial Vehicles (UAV) which are not able to overfly potential operational areas prior to the beginning of hostilities and are susceptible to anti-aircraft fire once hostilities begin. Second, while these sensors provide excellent data, some important terrain data themes (for example, soil types and soil moisture) remain very difficult to collect from remote sensors. Third, even the data from these sensors is uncertain. Users who are seduced by the high resolution and much higher accuracy of these data sets may trust them too completely, without realizing that some terrain effects are so sensitive to terrain values that even very small terrain data errors can cause large variations in predicted terrain effects.

In fact the availability of patches of high resolution, high accuracy data will contribute to the heterogeneous nature of the operational database and increase the challenge of managing and understanding the mix of data resolutions and qualities.



3.1 Representing Terrain Analysis Models as Probabilistic Models

The CCM product of figure 6, was produced using a Bayesian Network which implemented a CCM algorithm as a probabilistic model.

A Bayesian Network consists of a graphical model that contains nodes that represent uncertain variables and arcs that represent the qualitative influences between them. Local probability distributions on each node represent the strength of the influences. The uncertain variables can be continuous variables or discrete variables. Bayesian Networks can encode the type of information found in a logic system rule base, but are much more flexible in capturing relationships between related variables and uncertainty in the relationships between variables. Conditional probability distributions simplify the collection and exploitation of the knowledge required to develop the Bayesian Network model. Propagation algorithms determine the prior distribution of any variable in the network. When evidence on one or more of the uncertain variables is available, the propagation algorithm uses the evidence to propagate the remaining uncertainty throughout the network and determines the probability distribution of all other variables in a way consistent with the evidence.

A simple Bayesian Network is presented in Figure 7. This network is a simplified prototype of one that is used to propagate uncertainty in geospatial data through a CCM algorithm. Based on terrain data (slope, soil type, soil moisture, and vegetation) the CCM algorithm predicts the speed that a specific vehicle will be able to move cross country (off road). Existing algorithms produce a point estimate with no estimate of the prediction uncertainty.

As discussed above, the Bayesian Network is a graphical model, with nodes and arcs. The nodes represent uncertain variables; in this case they represent terrain variables, and the CCM speed. Each node can exist in one of a number of mutually exclusive states. For example, the nodes that represent vegetation type would have states that correspond to the vegetation classes in the database. Notice that the top row of nodes contains uncertain variables that represent the information in the database at a specific point on the ground (at one pixel). These variables are uncertain - at least until we read the database. The second row of variables represent a different set of uncertain variables, these represent the ground truth - which is unknown. The arcs between nodes represent the knowledge that there is a relationship - that will be defined as a conditional probability table - between the variables. For example, the arc between a database variable and the true terrain variable represents the knowledge that if we know the database value, then we have some information about the value of the true terrain.

The conditional probability distribution for each node is a table that defines the probability that the node will be in a particular state, given the value of its parents. For the example of the conditional probability for a true variable given a database variable, the conditional probability table can be derived directly from the error matrix or "confusion matrix" that defines the accuracy of the classification. In this simple Bayesian Network, this data quality information is fixed in the local probability distribution of the true variable node. In a more general Bayesian Network, more complex data quality models will represent the different data quality for different terrain data products. These more complex data quality models are discussed in section 5 below.

The other nodes and links in the Bayesian Network represent additional information about the problem domain. The soil strength node has two parents: true soil type and true soil moisture.

The conditional probability table for this node will define the probability distribution for the true soil strength given that we know the true soil type and the true soil moisture. The CCM speed node, has three parents, that reflect the knowledge that if we know the true vegetation type, slope, and soil strength, we can estimate the distribution of true CCM speeds. These probability tables encode the results of a particular CCM algorithm, modified to reflect the uncertainty in the model.




Figure 7. Prototype Bayesian Network for propagation uncertainty through a GIS Model.

To use this network, the terrain data is accessed, and the database nodes are instantiated to the terrain values in the database. The Bayesian Network propagation algorithm will then update the probability distribution for all other nodes in the network, and the result is a histogram of potential CCM speeds, showing their probability of occurrence, given the available terrain data.

Figure 8 shows the actual BN that was used to generate the probabilistic CCM product of Figure 6. This more complex BN encodes the ETL CCM algorithm (Pearson and Wright 1980).





Figure 8. Bayesian Network the encodes the ETL CCM algorithm.

Terrain analysis products are used in the military decision process as part of analysis of enemy or friendly COAs. Their use can be illustrated with a simplified decision model. In the situation shown in Fig 5 above, the blue commander considered the necessity to defend a potential Red advance through area B, which is assessed to be NOGO terrain. Figure 9 shows a decision model for this situation, for the case where Blue decision makers are not aware of the uncertainty in the terrain assessments.



The numbered items below are keyed to the numbers in the figure.

1) The terrain analysis assessment reports that the terrain is NOGO and therefore does not support the potential RED COA. The actual terrain analysis was performed using available data, which is subject to data quality and data currency problems, but the results reported to the commander do not include any assessment of uncertainty.


2) The Blue decision makers are not aware of the uncertainty in the terrain suitability assessment.

3) The Blue decision is based on the terrain assessment and does not consider the possibility that the terrain actually does support the potential Red COA. Therefore Blue is unlikely to defend against this Red COA.

4) There is a considerable probability that Red can achieve tactical surprise.

This situation, where one commander believed that the terrain did not support a specific enemy COA and his opponent was able to successfully carry out a surprise maneuver to achieve victory, has occurred over and over again throughout military history.

Figure 10 shows the same situation when the Blue decision makers are aware of the uncertainty of the terrain analysis assessment.



1) The terrain analysis assessment reports that the terrain is NOGO and therefore does not support the potential RED COA.

2) In this case, the Blue commander is aware of the uncertainty in the terrain assessment.

3) The Blue decision, whether to defend the COA, includes an understanding of the potential for uncertainty in the terrain assessment, and as a result is likely to defend against the COA.

4) As a result, Red cannot achieve a surprise by exploiting this COA.

Figures 9 and 10 are simplified representations of the kind of analysis used by Blue and Red commanders, but they do illustrate the central importance of understanding the uncertainty in terrain analysis assessment. When the Blue commander is aware of the uncertainty, he can defend against a dangerous Red COA. Not represented in Figure 10, the Blue commander has additional opportunities: for example to conduct a terrain reconnaissance, or to collect additional terrain data to reduce the uncertainty in the terrain assessment. With an appreciation of the uncertainty, the Blue commander has additional options that are not obvious when the uncertainty is ignored. This analysis has been from the point of view of a Blue commander assessing Red COAs. The same kind of analysis can be used by Blue in evaluating potential Blue COAs to identify potentials to achieve tactical surprise against Red.

In current practice, terrain analysis and COA analysis are performed by experienced human experts using automated tools. Human experts may provide the healthy skepticism necessary to recognize the potential risks in this kind of situation. In the near future, more and more of the analysis will be performed by automated algorithms in MPRS or command and control systems, and there will be fewer opportunities for human intervention.

3.1.1 Quiddity*Suite

The models used to implement the RKF TKB were implemented using QUIDITY*Suite, IET's software tools for Bayesian Inferencing. QUIDITY*Suite has been developed from the ground up over the past several years to design, analyze, simulate and refine systems that must work in situations with inherent uncertainty where incorrect results pose a very high risk, e.g., failure to

identify a potential enemy high speed avenue of approach. QUIDDITY*Suite consists of the following components:

QUIDDITY*Modeling – Rapidly model real-life situations using object oriented representations that combine and embody domain expertise, actual operating experience and performance data. Multiple inheritance, aggregation and abstraction are supported. QUIDDITY*Modeling is a knowledge representation language based on frames (a widely used knowledge representation in artificial intelligence) augmented to express uncertainties. In addition to frame (class) abstractions organized by "is-a" hierarchies inherited from the frame system, QUIDDITY*Modeling supports mechanisms to express uncertainties about the value of variables, the reference to instances, the existence of instances, and the type of instances. QUIDDITY*Modeling allows for expressing domain knowledge as pieces of Bayesian Networks, called BNFrags, in a modular and compact way, facilitating reuse. Instances of probabilistic frames are created dynamically, allowing situation specific probabilistic inference. The probabilistic inference is done by QUIDDITY*Inference using a Bayesian Network created dynamically from the current set of probabilistic frame instances. This generation of Bayesian Networks from QUIDDITY*Modeling utilizes QUIDDITY*Inference's local expressions to exploit all types of independence relationships to speed up the inference. QUIDDITY*Modeling is fully integrated with QUIDDITY*Script, allowing the user to define frames, create instances, and make situation-specific queries interactively.

QUIDDITY*Inference – Runs models to reveal the causes of observed effects about past and real-time performance plus produce ranked decision options for defending threats and capitalizing on opportunities. QUIDDITY*Inference is based on Symbolic Probabilistic Inference (SPI), one of only two known general solution algorithms for Bayesian Networks. In contrast to the alternate "join tree" approach to inference in Bayesian Networks, SPI has the following two important characteristics. First, SPI is query based. SPI extracts the minimum subset of a Bayesian Network that is necessary for each query, minimizing the amount of computation required for answering the query. In other words, the same query can be repeated many times from different points within the area of interest. Second, SPI has local expressions, an extension of Bayesian Networks, used to express local structure within a node. Local expressions can be used to instantiate many independence relationships including independence of causal influences and context-specific independence. SPI exploits these independence relationships in addition to the conditional independences inherent in Bayesian Networks for efficient inference in large Bayesian Networks. SPI has successfully computed queries for large "bench mark" Bayesian Networks, which the other inference algorithm is unable to compute. In addition, SPI's query-oriented approach allows for compilation of any probabilistic query into an efficient and small procedural code. In fact, because both the memory and CPU requirement of this generated code is fixed; it is readily usable in an embedded and/or real-time environment.

QUIDDITY*Script – Java-based command and scripting language for building, testing and production execution of models. IET's QUIDDITY*Script is an object-oriented scripting language designed specifically for Bayesian Network applications. Model builders can use it to dynamically construct Bayesian Networks from pre-built BNFrags, make situation-specific queries, and define and replace software components on the fly. In addition, the QUIDDITY*Script language can either be run interactively from a command line or via an API from within a larger software system - allowing automated control over construction and manipulation of Bayesian Networks

IET's Quiddity*Suite has been successfully used in a prototype geospatial application for management of uncertainty in geospatial data.^{4,5}

⁴ Wright, E. J., "Probabilistic Models In Geographic Information Systems – Bayesian Networks For Management Of Uncertainty", *Proceedings of the 4th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences*, Amsterdam, July 2000b.

⁵ Wright, E. J., "Understanding and Managing Uncertainty in Geospatial Data for Tactical Decision Aids", Ph.D. Dissertation, George Mason University, August 2002.

4. Terrain Suitability Knowledge Base

The TSKB contains knowledge about terrain factors that impact military operations. Figure 11 shows the components of the system.

1) Cyc and Cyc Knowledge Base. The TSKB includes an interface to Open Cyc and preexisting military knowledge in a Cyc knowledge base. The preexisting knowledge includes MicroTheories that contain knowledge about vehicles and military units, their composition, and the kinds of activities that they engage in. As necessary, additional logical terrain knowledge will be encoded in Cyc as part of a Cyc Terrain Knowledge Base.

2) GIS system. The TSKB will also include an interface to ERDAS Imagine, a commercial GIS system that is part of the military's deployed DTSS system. ERDAS provides capabilities to ingest, store, and manipulate spatial data, and the capability to define GIS models that perform standard transformations and manipulations of spatial data. The TSKB interface to ERDAS will provide a capability to access detailed spatial data necessary for reasoning about terrain suitability.

3) External Data Stores. The TSKB will have access to terrain and military data available in a wide range of external data stores. This includes terrain data in a variety of formats as well as other data important to military reasoning. Access to the data will be provided by OpenCyc's capability to interact with external databases and by ERDAS's capability to interact with external terrain databases.

4) Quiddity*Suite. The probabilistic reasoning engine for TSKB is provided by IET's Quiddity*Suite. Quiddity*Suite provides an efficient probabilistic inferencing engine, powerful object oriented probabilistic modeling tools, and a powerful scripting language for constructing and manipulating probabilistic models.

5) Bayesian Network Fragment Repository. The probabilistic models necessary for terrain suitability reasoning will be stored as Bayesian Network Fragments in a repository where they are available when needed to support terrain suitability reasoning.

6) TSKB custom software. The TSKB demonstration capability will include custom software that implements the interfaces to the other components and ties the entire system together. Some of the software will be written in Quiddity*Script, while the rest of it will be written in Java.

7) User Interface. The final component of TSKB is a user interface. This will provide an engineering interface to TSKB with sufficient functionality to exercise and demonstrate the functions of the TSKB.




Figure 11. The components of the Terrain Suitability Knowledge Base.

4.1 Logical Reasoning

In addition to probabilistic reasoning about terrain suitability, IET's TSKB will also exploit first order logic (FOL) systems. Integrating logical reasoning offers the opportunity to directly exploit the large volume of rule based (logical) information about military units, their activities, and their composition, as well as existing logical rules about factors for evaluation of military courses of action. For RKF Y3, IET developed a knowledge base containing representations for:

- Event types (Ambush, Convoy, Bivouac)
- Unit types (Tank, Mechanized Infantry, and Logistics Platoons and Companies, Guerilla Forces)
- Equipment types (Tanks, Humvees, Infantry Fighting Vehicles, Heavy and Light Trucks)
- Weapon types (Rifles, Machine Guns, Handguns, Tank Guns)
- Typical vehicles and weapons for each unit type
- Application of suitability factors and suitability models (CCM, LOS, Cover)

Such information can be represented directly in a probabilistic network, where a probability of 1 is used to represent True, and 0 for False. However, FOL reasoning of the sort used to represent such facts as "Tank Platoons have X number of tanks assigned" does not require the same reasoning framework as that used to handle uncertainty. It is more efficient to exploit knowledge bases, well suited to first-order logic reasoning, both as a repository for such knowledge and as the inference engine used.




Figure 12 Overview of First Order Logic system use

Using FOL, we can leverage a small amount of user input, shown in the blue box in Figure 12, in order to identify which suitability factors, vehicles and weapons to use in calculations for cross-country and on-road mobility, and line of sight (LOS). Determination of such factors is critical for COA analysis. In addition to identification of relevant data, FOL systems can be used as

storehouses for data. For instance, once we have determined that a unit has a M1-A1 Tank as one of its vehicles by querying the FOL system, we can also retrieve the M1-A1's vehicle profiles from the FOL system and use those values in our CCM model. Using a FOL system with a knowledge base in this manner makes it relatively easy to add vehicles or units, and change values associated with such entities. All changes are made in one central knowledge base and the uncertainty reasoning queries for relevant values as necessary. IET used FOL systems in combination with IET's uncertainty reasoning capabilities to achieve the most efficient division of labor and the best overall results.

Integrating logical reasoning with IET's approach to uncertainty also provides opportunities to exploit some of the capabilities currently being developed by other IET research projects.

4.2 Logic Models

For RKF Y3, we have developed two main model types—Unit models and Activity models. In addition, we have researched Vegetation models, using work done on geographic analogues as a starting point. Each model type is described below.

4.2.1 Unit Models

In evaluating COAs, one must consider the units involved. Our interest in units is currently restricted to the vehicles and weaponry associated with each unit. The term unit is used rather loosely, capturing conventional troops such as Mechanized Infantry Platoons as well as less traditionally organized troops such as Guerilla Forces. For each unit, we specify the vehicles attached to the unit. If the vehicles have weapons included (such as a tank with an integral gun), we represent the weapon as part of the unit's available weaponry. In addition, we list the usual weapon associated with troops attached to each unit type. This is relevant because not all troops carry the same weaponry—infantry would be expected to carry much heavier weaponry than would logistics troops. This becomes important when we are trying to determine likely ambush scenarios—an enemy would most likely attack more heavily or from better-guarded positions if attacking heavily armed troops versus lightly armed troops.

In future work, our activity models will become richer, including additional features such as training level, expected tactics, ground troop movement, and unit intentions.

4.2.2 Activity Models


Activity models are used to determine the most relevant equipment and weapons to use when determining suitability. For each unit, the system can identify the attached vehicles and the expected weapons (both for the vehicles and for the troops). For a given activity, particular vehicles and/or weapons are used when calculating suitability factors. For example, when determining an ambush site, the ambush location(s) need to place the ambushing troops within range of the ambush targets. To determine range, we currently identify the most likely weapon to be used in the ambush for the identified troops and then use that weapon's range in our calculations. This information is hard-coded in the system for both weapons and for vehicles, which are used for mobility models. In the future, this information could be determined via rules written in the FOL knowledge base. An example we have discussed is selecting the vehicle to use for CMM calculations based on a variety of features such as maximum road speed, weaponry, or visibility.

Currently activity models are limited to Convoy, Bivouac, and Ambush. In future work, we would expand the known set of activities and continue to improve the reasoning used to determine relevant weapons and vehicles for each activity type.

4.3 Probabilistic Models

4.3.1 Data Quality Models

This sub section describes data quality models that provide models of feature classification accuracy. Figure 13 shows a basic error model for vegetation class accuracy. The model is a Bayesian Network with a node to represent the true vegetation class and another to represent the vegetation class in the simulated terrain data. The Bayesian Network on the left has one additional node to represent the error. This simple model can be extended (as shown on the right) to a model where the error depends on the spatial and spectral resolution of the source used to create a dataset.







Figure 14 is a more complex model that includes the relationships between several terrain features and a common error model.



work done previously in order to correctly identify a sub-region with its larger, representative biome and then matching the identified biome with previously collected vegetation data.




Figure 17. Example of inferred vegetation data for terrain reasoning.

In addition, we have developed a small set of models for inferring missing vegetation data if we have enough supporting data known. For example, in Figure 17, we show a possible path for inferring Concealment from a variety of vegetation and vegetation-related factors. If we know Stem Diameter, we can determine Concealment directly. Lacking Stem Diameter, we may also use, walking up the chain, Vegetation Cover, or Average Annual Rainfall and Soil Type, to infer Concealment. Each inference step introduces more uncertainty into the final determination. A benefit of our approach to uncertainty is that included in the final determination will be an indication of the level of uncertainty attached to the system-generated answer.

Our work regarding vegetation models generated several ideas for future terrain reasoning research. However, given the time and funding limits, the majority of this work has not been implemented.

4.3.2.1 Geographic Relationships

Bayesian Networks can be used to represent the relationships between geographic features represented in terrain data products. Some examples were presented in section 2 (Figures 4, and 5).






Figure 19 shows a more complex relationship model that includes additional geographic features and attributes. This model includes relationships between geology, climate, landform, elevation, land use, and vegetation. This model (and others like it) would be used to infer needed terrain data values from values that are available in available data sets.

Figure 20 shows an extension to this model as the result of inclusion of a new vegetation attribute – vegetation understory. This new attribute can also be used to illustrate the need for inferring terrain data values. Research by US Army Topographic Engineering Center (TEC) and White Sands Missile Range (Krause, *et. al.*, 2001) has identified the importance of detailed information about vegetation understory in forested areas (height and density of low branches and bushes beneath the canopy). Understory has a dramatic effect on line of sight and models of engagement ranges used in modeling and simulations. Unfortunately, understory data is not available in any existing terrain data products. A probabilistic model like this provides a means to infer understory attributes from other available data.



4.3.3 Terrain Assessment Models


Two terrain assessment models were implemented as probabilistic models in the TSKB. The



first was the CCM Tactical Decision Aid which was illustrated in Figure 8, above. For the TSKB, the CCM model was implemented as Bayesian Network Fragments using Quiddity*Modeler. Several of these frames are discussed in Section 5.4 below. Figure 21 is an example of a probabilistic CCM product produced from the Bayesian Network model.

The second suitability model implemented was probabilistic LOS, which considers the elevation errors in a Digital Elevation Model (DEM) and computes a probability of line of Sight existing between two points.

It is clear that LOS predictions will be affected by elevation errors in the DEM. DEM accuracies are often specified with an absolute elevation error probability (i.e. linear error 90%). An absolute elevation error is the total error in each DEM elevation. It includes the vertical bias and the random error. However, a vertical bias will have no impact on LOS predictions, which depend only on the relative elevation errors at the observer, the target, and any obstruction.



The LOS model used in this study is shown in Figure 22. The values $e1$, $e2$, and $e3$, are the elevations from the DEM of the observer, a potential obstruction, and the top of the target. The values $d2$ and $d3$ are the distances from the observer to the potential obstruction, and to the target. This model assumes a straight line LOS, no Earth curvature or refraction, and no obstructions from vegetation or built up areas. These restrictions could be lifted in a more complete application.

LOS is calculated with the following algorithm:

$$H_{LOS} = ((e1 - e2)/d2)d3 - (e3 - e1) \quad (42)$$

If $(H_{LOS} > 0)$ then it passes above the target and LOS is blocked, else LOS is clear, where H_{LOS} is the Height of the line of sight (that intersects the obstruction) above (or below) the top of the target.

To calculate the accuracy of the LOS prediction, we used the standard error propagation formula from statistics (Mikhail 1976):

$$\begin{aligned} \mathbf{Y} &= \mathbf{F}(\mathbf{X}) \\ \Sigma_{yy} &= \mathbf{G} \Sigma_{xx} \mathbf{G}^T \\ \mathbf{G} &= \delta \mathbf{F}(\mathbf{X}) / \delta \mathbf{X} \end{aligned} \quad (43)$$


In the LOS case, the independent variables ($\mathbf{X}$) are $e1$, $e2$, and $e3$. The dependent variable ($\mathbf{Y}$) is H_{LOS} . The function $\mathbf{F}(\mathbf{X})$ is the equation for H_{LOS} above. To compute the accuracy of H_{LOS} (Σ_{yy}),

we need the variance-covariance matrix of $\mathbf{X}$, (Σ_{xx}) and the matrix $\mathbf{G}$, which is the partial derivatives of the function with respect to the elevations ($e1, e2, e2$):

$$\mathbf{G} = [1 - d3/d2, d3/d2, -1] \quad (43)$$

The result of the above calculation is the matrix Σ_{yy} . In this case it is 1 x 1 matrix (a scalar), which is the variance of H_{LOS} . Assuming a Normal distribution, the probability that LOS exists can be estimated as shown in Figure 23.

The final step is to determine the variance-covariance matrix of $\mathbf{X}$, (Σ_{xx}). The diagonal elements are the variances of the elevation estimates, the off diagonal elements are the covariances



between the different elevation estimates.


Assuming that the elevation errors are spatially correlated, as a function of distance, the covariance between two elevations will be a function of the distance between the points.

If a suitable test DEM is available, an empirical covariance function can be estimated from the data for use in constructing the variance covariance matrix required for LOS error propagation.

This algorithm was coded in a Java application that works with standard digital elevation models from NGA.

4.4 Knowledge-Based Model Construction Implementation

In Year 3, IET has implemented a proof-of-concept knowledge base to facilitate the construction of situation specific probabilistic networks relevant to COA evaluation. Mahoney and Laskey (Mahoney and Laskey, 1998), describe Knowledge-Based Model Construction (KBMC) as the



process of constructing a model for a problem instance from a knowledge base representing generic domain entities and their interrelationships. Such a system should include a knowledge base of reasoning rules and network fragments, search operators for retrieving problem-relevant knowledge base elements, network construction operators, network evaluation operators, and model construction control mechanisms. Objectives for a KBMC system are to minimize costs of representation, retrieval, construction and evaluation, while providing accurate responses to queries. In particular, our knowledge base should allow us to query for reasoning paths from known or knowable data to properties of interest so as to facilitate the construction of efficient situation specific networks. Mahoney and Laskey note in 0:

In more complex domains it is necessary to reason about a variable number of entities that may be related to each other in varied ways. It is also necessary to reason about and distinguish between multiple instances of a given complex pattern of entities and relationships. In such domains it is infeasible to construct a complete belief network encompassing all the situations one might encounter in problem solving.

During the Y3 RKF work, IET identified the implementation of RKF tools for KB construction and exploitation as a means by which the extant RKF tools could be reused for purposes of utilizing a knowledge base to facilitate dynamic network construction. In particular, this extends IET's KBMC capabilities by introducing search operators for retrieving problem relevant knowledge base elements. The RKF tools were largely focused on the elicitation of logic or rule-based knowledge, so it was this kind of knowledge that we attempted to implement.

The challenge which KBMC techniques address is as follows. Utilizing QM frames, or probabilistic relational models, results in the creation of a large number of frames. For example, in our terrain KB we had a number of distinct frames some of which are summarized below, i.e., Context, TrueTerrain, CCM and Vehicle frames. We have deleted some of the slots, and probability distribution information for brevity.

```

frame Context isa Frame
  slot season
    facet domain = [spring, summer, fall, winter]
    facet distribution = [.25,.25, .25, .25]
  slot lastRain
    facet domain = []
    facet parents = [season]
    facet distribution = function s {}
  slot topography
    facet domain = [plains, hilly, rugged, mountainous]
    facet distribution = [.2, .45, .25, .1]
end;

```

```

frame Vehicle isa Entity
  slot predictedCCM
    facet domain = CCM
  slot gradability
  slot maxRdSpeed = 75.0
  slot width
  slot overRideDiam
  slot vci1
  slot vci50
  slot maxFordDepth
  slot maxStreamVel
  slot vehAppAngle
  slot maxVertObs
  slot DitchCrossCap
  slot speed
    facet domain = Continuous
    facet partition = [0, .5, 3.0, 6.0, 10.0, 15.0, 20.0, 25.0,
                      30.0, 40.0, 50.0, 65.0, 100.0]
    facet parents = [maxRdSpeed]
    facet distribution = function mrs { }
  slot obsSpeed
    facet domain = Continuous
    facet parents = [speed]
    facet partition = [0, .5, 3.0, 6.0, 10.0, 15.0, 20.0, 25.0,
                      30.0, 40.0, 50.0, 65.0, 100.0]
    facet distribution = function sp { }
  slot speedConstraint
    facet domain = [true, false]
    facet parents = [exists,predictedCCM.ccm,speed]
    facet distribution = function ex, ccm, sp { }
    facet value = true
end;

```

Figure 25 Context and Vehicle Frames

Note that these frames (Figure 25 and Figure 26) indicate qualitative relationships between properties of different kinds of objects in the frame system. If we want to reason about CCM property f4, then the frames indicate that we need to know the terrain and various properties of the vehicle(s) that is/are relevant to reasoning about that property. Or if we have evidence about vehicle speed, etc., it is useful to know which further properties we can reason about. Hence implementing a BN to reason about CCM property f4 may require constructing or selecting instances of a number of different frames and defining or ensuring that they are in an appropriate relationship with one another. For example, that slot f2 on CCM is influenced by various properties of the object that fills the vehicle slot and properties of the value that fills the terrain slot, the properties of the terrain are further influenced by the value of context that fills the ‘context’ slot for the relevant value of a TrueTerrain instance. Hence, in a reasoning context in which we are interested in the value of f2 we need to know which frames to instantiate and/or which objects are of interest. It will be important to have some instance of the frame CCM but also some instance of Vehicle, TrueTerrain and Context. Furthermore, not any instance of these frames are salient, we need to implement the ones that bear the appropriate relationships, e.g., an instance of Vehicle that fills the ‘vehicle’ slot in the instance of CCM about which the user is reasoning.

In more complex domains it is necessary to reason about a variable number of entities that may be related to each other in varied ways. It is also necessary to reason about and distinguish between multiple instances of a given complex pattern of entities and relationships. In such

domains it is infeasible to construct a complete belief network encompassing all the situations one might encounter in problem solving.⁷

```

frame TrueTerrain isa Frame
  slot context
    facet domain = Context
  slot slope
    facet domain = [A,B,C,D,E,F,G]
    facet parents = [context.topography]
    facet distribution = function t { }
  slot veg
    facet domain = []
    facet distribution = [.5, .2, .1, .1, .07, .03]
  slot vegP
    facet domain = [yes, no]
    facet parents = [veg]
    facet distribution = function v { }
  slot vegSD
    facet domain = Continuous
    facet parents = [veg]
    facet distribution = function v { }
  slot vegSS # units: meters
    facet domain = Continuous
    facet parents = [vegSD,veg]
    facet distribution = function sd, v { }
  slot surf
    facet domain = [NoEffect, Stony, StonyRkOc, Karst,
                    Quarry]
    facet distribution = [.7, .1, .1, .07, .03]
  slot surfR
    facet domain = Continuous
    facet parents = [surf]
    facet distribution = function s { }
  slot soilType
    facet domain = []
    facet distribution = []
  slot soilMoisture
    facet domain = [wet,moist,dry]
    facet parents = [context.season,context.lastRain]
    facet distribution = function s, lr {switch s { } }
  slot soilRCI
    facet domain = Continuous
    facet partition = []
    facet parents = [soilType,soilMoisture]
    facet distribution = function st, sm {switch sm { } }
end;

```

```

frame CCM isa Frame
  slot vehicle
    facet domain = Vehicle
  slot terrain
    facet domain = TrueTerrain
  slot s1
    facet domain = Continuous
    facet parents = [vehicle.maxRdSpeed,
                    vehicle.gradability,terrain.slope]
    facet distribution = function m, g, s { }
  slot f1
    facet domain = Continuous
    facet parents = [vehicle.width,terrain.vegSS,
                    terrain.vegSD]
    facet distribution = function vw, ss, sd { }
  slot f2
    facet domain = Continuous
    facet parents = [vehicle.width,vehicle.overRideDiam,
                    terrain.vegSS,terrain.vegSD]
    facet distribution = function vw, vod,ss, sd { }
  slot s2
    facet domain = Continuous
    facet parents = [terrain.vegP,s1,f1,f2]
    facet distribution = function vegP, s1, f1, f2 { }
  slot s3
    facet domain = Continuous
    facet parents = [s2,terrain.surfR]
    facet distribution = function s2, sr { }
  slot f4
    facet domain = Continuous
    facet parents = [terrain.soilRCI,vehicle.vci1,
                    vehicle.vci50]
    facet distribution = function rci, vci1, vci50 { }
  slot ccm
    facet domain = Continuous
    facet parents = [s3,f4]
    facet partition = [0, 0.5, 3.0, 10, 20, 30, 45, 100]
    facet distribution = function s3, f4 { }
end;

```

Figure 26 True Terrain and Cross Country Mobility Frames

```

slot speed
  facet domain = Continuous
  facet partition =
    [0, .5, 3.0, 6.0, 10.0, 15.0, 20.0, 25.0, 30.0, 40.0,
     50.0, 65.0, 100.0]
  facet parents = [maxRdSpeed]
  facet distribution = function mrs { }
slot obsSpeed
  facet domain = Continuous
  facet parents = [speed]
  facet partition =
    [0, .5, 3.0, 6.0, 10.0, 15.0, 20.0, 25.0, 30.0, 40.0,
     50.0, 65.0, 100.0]
  facet distribution = function sp { }

```

Figure 27 Example of parent-child slot relationships

The challenge that we face is to be able to reuse these frames and instantiate and integrate them according to the reasoning task at hand, and the nature of the evidence that is available. In order to effect this, we model the qualitative relationships represented in the various frames into a large graph containing information about all frames in our system using a Perl script to translate the frames into a KB representation. If a property on a frame or object is affected by another property of that

object, we represent each property as a node and then create an edge between them. For example (Figure 12), on the vehicle slot, speed, we see that a parent node of that slot is ‘maxRdSpeed’.

We represent this in our graph as: `edge(vehicle*maxRdSpeed,vehicle*speed)`.




Figure 28: Graphical representation of links between frame slots

However, if we posit a link between a slot in one frame to a slot in another we require more information to maintain accuracy and even structure. We need to know what the relation is between the two frames that is presupposed by the Bayesian qualitative link. So, for example, consider the following CCM frame:

```

frame CCM isa Frame
  slot vehicle
    facet domain = Vehicle
  slot terrain
    facet domain = TrueTerrain
  slot s1
    facet domain = Continuous
    facet parents = [vehicle.maxRdSpeed,
                    vehicle.gradability,terrain.slope]
    facet distribution = function m, g, s {}
  slot f1
    facet domain = Continuous
    facet parents = [vehicle.width,terrain.vegSS,
                    terrain.vegSD]
    facet distribution = function vw, ss, sd { }
  slot f2
    facet domain = Continuous
    facet parents = [vehicle.width,vehicle.overRideDiam,
                    terrain.vegSS,terrain.vegSD]
    facet distribution = function vw, vod,ss, sd { }
  slot s2
    facet domain = Continuous
    facet parents = [terrain.vegP,s1,f1,f2]
    facet distribution = function vegP, s1, f1, f2 {}
end;

```

Figure 29. CCM frame example

There are causal links between CCM.s1 and Vehicle.gradability but to make this meaningful we require the extra information contained in JPF/QM, i.e., the fact that the causal relation is between the value of S1 of CCM and the gradability of a vehicle (i.e., the value of the slot 'vehicle') and the slope of the terrain (i.e., the value of the slot 'terrain').

Hence, it would not suffice to simply assert that there is a causal relationship between CCM.s1 and Vehicle.gradability, i.e. `edge (Vehicle.gradability CCM.s1)`. If we want to reason over actual CCM instances we need to ensure that we are considering the relevant instance of Vehicle and terrain should more than one exist, i.e., we want the instance of Vehicle salient to the particular CCM model we are implementing.

It is with an eye to retaining this information, which is centrally relevant to indexing and reconstructing that we determine whether a slot references another frame. If it does, we make the link between two properties indirect introducing an intervening node that specifies the nature of the relationship between the distinct frames that are connected.

So, rather than asserting the simple relation we assert:

```
edge(Vehicle.gradability [Vehicle.gradability,vehicle])
edge([Vehicle.gradability,vehicle] CCM.s1).
```

The semantics of the resulting graphs that we implemented can be described as follows. Suppose there is a path from FrameA.slot1 to FrameB.slot2 with an intervening node '[FrameA.slot1, rel]', i.e., `path(... FrameA.slot1, [FrameA.slot1, rel], FrameB.slot2...)`. This means that in an instance of FrameB the value of slot2 is influenced by the value of slot1 of a particular instance of FrameA, i.e., the instance that bears the relation 'rel' to the instance of FrameB. See Figure 30 for an example of how we add nodes to a graph to show how distinct frames salient to the reasoning must be related. On the left of the figure there are a number frames and slots presented as nodes. If the s1 slot of a CCM frame instance is influenced by the slope slot of a TrueTerrain frame instance representing the terrain of that CCM, then we represent the requisite relationship, i.e., between the two frame instances in an intervening node as shown on the right, i.e., we show that the influence of the TrueTerrain slope slot is predicated on the assumption that the TrueTerrain frame bears the 'slope' relation.




Figure 30. Representing relations between frames requisite for causal influence

So, in the case where the relevant causal relationship exists between two different frames we introduce new nodes that aid in assembling frames for a situation specific BN. A partial representation of the graph representation of causal relations relevant to our Y3 CCM model is given below:

```

edge(vehicle*maxRdSpeed,vehicle*speed).
edge(vehicle*speed,vehicle*obsSpeed).
edge(vehicle*exists,vehicle*speedConstraint).
edge(cCM*ccm,[cCM*ccm,predictedCCM]).
edge([cCM*ccm,predictedCCM],vehicle*speedConstraint).
edge(vehicle*speed,vehicle*speedConstraint).
edge(context*season,context*lastRain).
edge(context*topography,[context*topography,context]).
edge([context*topography,context],trueTerrain*slope).
edge(trueTerrain*veg,trueTerrain*vegP).
edge(trueTerrain*veg,trueTerrain*vegSD).
edge(trueTerrain*vegSD,trueTerrain*vegSS).

```

Part of our Y3 effort involved creating code to generate appropriate graph specifications, as described above, implementing a set of frame descriptions as input. The code for performing this is available from IET upon request. Depending upon the kind of evidence that we have at hand, the properties in which we are most interest, etc., we need to determine which frames can be instantiated and the relationships that exist between them. Using graph traversal algorithms we can then traverse the graph to identify the relevant causal influences for purposes of constructing a BN to perform the reasoning. We implement this in prolog and show the query results below:

```

5 ?- path(X,cCM*s2,[X],Path).
X = cCM*s2
Path = [cCM*s2] ;
X = context*topography
Path = [context*topography, [context*topography, context], trueTerrain*slope,
[trueTerrain*slope, terrain], cCM*s1, cCM*s2]

```

We can further analyze these results to build a more comprehensive network as necessary. For example, we might query for reasoning paths to ‘trueTerrain*slope’ to determine whether other frames or evidence should be incorporated to facilitate reasoning about that property.

In terms of JSPI implementation we interpret the results as follows. First, we parse out all the frame names, context, trueTerrain, cCM, and therefore instantiate frames as follows:

```

cntx = Context->makeInstance();      #create an instance of the Context
                                     frame
terr = TrueTerrain->makeInstance();   #create an instance of the
                                     TrueTerrainFrame
ccm1 = CCM->makeInstance();           #create an instance of the CCM frame

```

We then use the information in the path to ensure that the instantiated frames are in the appropriate relation, i.e., that path (context*topography, [context*topography, context], trueTerrain*slope) means that:

terr->context = cntx; #ensure that (context terr cntx)

and path(trueTerrain*slope, [trueTerrain*slope, terrain], cCM*s1)

ccm1->terrain=terr #ensure that (terrain ccm1 terr)

This produces frames as needed and ensures that the frame instances are in the appropriate relation. As a result, users unfamiliar with the KB of fragments are able to query for relevant reasoning paths and build appropriate situation specific networks. Further, this implements tools necessary or useful in the creation of a system in which situation specific networks can be created automatically.

4.5 Scenarios

To demonstrate the application of the TSKB, IET has developed a set of scenarios based on predicting suitability, from the enemy's point of view, for sites where an enemy could ambush a friendly convoy. Suitability will be predicted based on terrain factors, but also based on enemy unit type, equipment, and potential objectives. A typical scenario involves the specification of Blue and Red unit types, Red forces' intent, identification of controlling agent for the territory under consideration, and activity considerations.


Blue Unit Type: 15 Trucks Red Unit Type: 20 guerilla fighters Red Intent: Capture supplies Control of Area: Open, Red-friendly Activity Considerations: ● Red will need to: ◆ move unit into area ◆ move and conceal unit at ambush site ◆ have vehicles to remove supplies (or leave trucks intact) ◆ have somewhere to hide vehicles and/or supplies to avoid counterattack	
--	---

Figure 31. Scenario Example: Ambush with intent to capture supplies

The generated scenarios have been used to drive the development of both Bayesian models and first-order logic reasoning. We have identified several terrain suitability factors associated with each scenario. Below we show factors associated with the scenario above.




Figure 32. Example of terrain suitability factors

Scenarios have driven the development of supporting knowledge for both our knowledge base work and suitability model development.

4.6 Application

Figure 33 illustrates the application of the TSKB as part of a larger automated predictive battlespace and COA analysis capability. A future predictive battlespace capability will generate a set of potential enemy (or friendly COAs) based on available evidence and the current assessment of the battlefield. Part of the predictive battlespace capability is a terrain engine that is capable of evaluating the suitability of the terrain for the alternative COAs. The TSKB is a prototype of the functionality required of this terrain engine. It is capable of determining which terrain factors are important for each COA, determining the necessary terrain features and attributes, the available terrain data - to include inferring mission values if necessary, calculating a probabilistic assessment of each factor, and combining them into an overall assessment of the terrain suitability for each COA.




Figure 33. The application of the TSKB as part of an automated COA evaluation capability.

5. Extension to Support FFD or Other Terrain Data Sets

The current implementation of the TSKB is able to work with two specific terrain analysis products: ITD, and VITD, which are the terrain data products that are most widely available today. However under NGA's current concept for terrain data support, future terrain data will be provided in two new terrain products: Foundation Feature Data (FFD) and Mission Specific Data Sets (MSDS).

FFD is intended to provide sufficient terrain information to support contingency planning and initial crisis response planning, while being inexpensive enough to produce that it can be made available for large regions of the world.

MSDS is a more detailed terrain data set that provides detailed information to support specific military operations. MSDS data is expensive to produce, and so will only be produced for specific areas where military commanders have identified a specific requirement. In addition, MSDS allows commanders to identify the specific terrain features and attributes needed for their specific operational mission.

The implication of these new products is that any automated terrain analysis process must be able to accept a variety of input data formats with a wide range of terrain data content. This section outlines the developments necessary to update the TSKB to accept this expanded range of terrain data products and content. There are two issues that these developments must address: 1) ability to exploit an expanded range of terrain data formats, and 2) the ability to exploit the different terrain data content available in new terrain data products.

5.1 Expanded Range of Terrain Data Formats


The TSKB is built around an interface with the ERDAS GIS software. This GIS package is part of the DTSS terrain analysis system used by the Army, and is also part of the standard software used by the Marine Corps and other services. ERDAS provides capabilities to import and export all standard DoD terrain product formats, as well as a very large set of civil government and commercial data formats. Because it is a standard part of military terrain analysis systems, ERDAS is being continuously updated with new import and export capabilities to support new or evolving terrain data formats. As a result, TSKB will have access to new terrain product formats as they become available.

5.2 Expanded Range of Terrain Data Content

The capability to accommodate the expanded range of data content available from FFD, MSDS or other terrain data products, is based on recognition that terrain suitability Bayesian Networks encodes two distinct types of information. This is illustrated in Figure 34. The top row on Bayesian Network nodes represent the database values of slope (DBSlope), vegetation stem spacing (DBVegSS), vegetation stem diameter (DBVegSD), ground roughness (DBGR), soil type (DBSoil), and soil moisture (DBSMoist). These nodes are connected to a second row of Bayesian Network nodes that represent the (unknown) true values for these terrain variables. All of the information required to construct this portion of the Bayesian Network can be supplied by the data quality information that accompanies each GIS data layer. If a new terrain product

contains data content (features and attributes) that was used to construct the Bayesian Network, then exploiting the new product requires only new data quality models that captures the accuracy of the new product. (Data quality models are discussed in Section 5.3.1).

The rest of the Bayesian Network is independent of the data. It can be used as is for any data set, of any quality. The information required to define this portion of the Bayesian Network is available directly from the TDA algorithm implemented; in this case the ETL CCM algorithm. This portion of the Bayesian Network represents the Joint Probability Distribution (JPD) of the true terrain variables and the CCM speed, that can be used to compute the conditional probability distribution for CCM speed given true terrain. In general, any TDA or terrain suitability algorithm could be implemented as a Bayesian Network that provides the conditional probability of the TDA result given the true terrain values. Note that any TDA implemented as a Bayesian Network could, and should, include model uncertainty that represents the inaccuracies of the TDA algorithm.



If, as is the case for FFD and most versions of MSDS, the new terrain data product does not contain all of the terrain features and attributes used by the TDA algorithm, or in the case that the new terrain product uses a different set of feature types and attributes to describe the terrain, then several options are available:

- Use inferred values for values missing from the product. In the worst case, the terrain suitability model can use a prior distribution for the missing values. Prior distributions can be defined by geographers, who are experts in the features and regions of interest. Alternatively

prior distributions can be learned from other available terrain datasets from similar geographic regions. The idea of geographic analogues, discussed in section 5.3 above, can be used to identify geographic experts or analogous datasets, even when no region specific experts or terrain datasets are available.

- Use geographic relationship models, as discussed in Section 5.3 above, to infer missing values from available terrain features and attributes. These relationship models can be constructed by geographers, who are experts in the features and regions of interest. The idea of geographic analogues, discussed in section 5.3 above, can be used to identify geographic experts even when no region specific experts are available.

- In some case, the same actual features may be described in the new terrain product, but features and attributes may be defined using a different ontology. The terrain suitability models implemented in the TSKB were constructed to work with the available ITD and VITD terrain data products, which use the FACC as a terrain ontology to describe features and attributes. The CCM algorithm was also designed to work with (an earlier version of) these data sets. Use of a different ontology will require transformation to the ontology in use by the TSKB. As more and more different types of terrain products become available, a larger and larger set of transformations will be required.

5.2.1 Interoperable Definition of Terrain Features and Attributes

The EDCS was discussed in Section 3.1.3 above. It was developed to provide an interoperable ontology for use in environmental data application originally for M&S, but now being extended to operational applications. The EDCS includes transformations between different terrain ontologies, for example between FACC and EDCS. Some of the transformations are exact, while some transformations are inexact and will introduce some uncertainty into the transformation. These inexact transformations could easily be represented as Bayesian Network Fragments. If the TSKB were modified so that the suitability models were defined using the EDCS ontology, that would facilitate the future use of additional terrain data products.

5.3 Tasks to Extend TSKB to Support FFD.

The following tasks would be needed to extend TSKB to use FFD.

Task 1. (optional) Modify the terrain suitability models to use the ECDS terrain ontology vs the current FACC based ontology. This task is not strictly required, but would provide foundation that would make it much easier to add additional terrain data products in the future.

Task 2. Data Quality Models. This task would develop the data quality models that represent the thematic accuracy of the feature and attribute data in the FFD. This task will involve some research, and knowledge elicitation, because the thematic accuracy of the FFD product is not defined in the FFD product specification. The data quality models would be represented as Bayesian Network fragments, for use in the TSKB.

Task 3. Geographic Relationship Models. This task would develop geographic relationship models that would allow inferencing for terrain features and attributes that are not contained in the FFD product. The geographic relationship models would be represented as Bayesian Network fragments, for use in the TSKB.

Task 4. Logic rules to support model construction. This task would develop the logic rules that would guide the construction of the specific Bayesian Network that will use the appropriate data quality and geographic relation ship models with the suitability models perform inference for terrain suitability.

6. Conclusions

IET's TSKB is a proof of concept prototype for a terrain engine for evaluating the suitability of terrain to support COAs.

The TSKB provides a integrated logic and probabilistic reasoning engine that uses logic to determine the terrain factors that are important, determine the required terrain features and attributes, as well as to construct the specific Bayesian Networks needed to assess terrain suitability. The TSKB uses probabilistic reasoning to asses the effect of different terrain factors in a way that takes into account the data quality of the terrain data, and assesses the uncertainty of the resulting terrain assessment.

The use of uncertainty in the terrain analysis process provides a mechanism that reveals potential risks - and opportunities, that would be missed if uncertainty is ignored.

The implementation of the TSKB was done in a way that can readily be extended to exploit new terrain data products.

7. References

- Birkel and Miller, 2000 "Developing a Common Terrain Data Model for the Army", ECDM, TEC TEM, briefing presented at the USATEC Technical Exchange Meeting, 26&27 Sep 2000.
- Castillo, E., Gutierrez, J., and Hadi, A., 1997, *Expert Systems and Propbabilistic Models*, Springer-Verlag, New York
- Congalton, and Green, 1999, *Assessing the Accuracy of Remotely Sensed Data: Principles and Practices*, Lewis Publishers.
- Cowell, R., Dawid, A., Lauritzen, S., and Spiegelhalter, D., 1999, *Probabilistic Networks and Expert Systems*, Springer-Verlag, New York
- Cressie, N.A.C, 1993, *Statistics for Spatial Data*, John Wiley & Sons, New York.
- Davis, T.J., 1998, "Modeling and Visualizing Multiple Spatial Uncertainties", <http://www.elsevier.nl/inca/homepage/miss/cageo/davis/article.htm>, Elsevier Science
- DMS, 1982, *Analysis of Natural Features*, Defense Mapping School (DMS) Student Text.
- FGDC, 1998, "Federal Geographic Data Committee Content Standard for Geospatial Metadata", FGDC-STD-001-1998.
- Forbus, K., Usher, J., and Chapman, V. 2003. "Sketching for Military Courses of Action Diagrams", Proceedings of the Intelligent User Interfaces Conference (IUI-03), January 2003. Miami, Florida.
- Gelman, A., Carlin, J., Stern, H. and Rubin, D., 1995, *Bayesian Data Analysis*, Chapman and Hall, London
- Gilks, C.J., Thomas, A., and Spiegelhalter, D., 1994, "A language and a program for complex Bayesian Modeling", *The Statistician* , 43, 169-78.
- Goodchild and Longley, 1999, *The Future Of GIS And Spatial Analysis, Geographical Information Systems, Volume 1, Principles and Technical Issues*, Longley, Goodchild, Maguire, and Rhind, editors, John Wiley and Sons, 1999.
- Goodchild, 1992, "Closing Report NCGIA Research Initiative 1: Accuracy Of Spatial Databases I", National Center for Geographic Information and Analysis, University of California.
- Heckerman, D., 1996, "Bayesian Networks for Knowledge Discovery", *Advances in Knowledge discovery and Data Mining*, Editors Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P., and Uthurusamy, R., The MIT Press, Cambridge, MA, 1996.
- Heuvelink and Lemmens, *Accuracy 2000*, Proceedings of the 4th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences, Amsterdam, July 2000.
- Heuvelink, G., 1998, *Error Propagation in Environmental Modelling*, Taylor & Francis, London.

- Holland, Wilson, Chung, Breckenridge, and Shaw, 2000, "An object-oriented approach to characterization of the dynamic littoral environment at Duck, NC", Submitted to Journal of Coastal Research, May 2000.
- ISO TC 211 Working Group 3, 1996, Geographic Information - Metadata (ISO 15046-15), Working Draft, March 29, 1996.
- Jensen, F.V., 1997, *An Introduction to Bayesian Networks*, Springer-Verlag, NY.
- Krause, and Clark, 1993, *Representing Uncertain Knowledge: An Artificial Intelligence Approach*, Kluwer Academic Publishers, Oxford.
- Krause, Campbell, and Puffenberger, 2001, "Estimating Understory Stem Density from Overstory Structural Characteristics", US Army Engineer Research and Development Center (ERDC/TEC) publication TR-01-1, Alexandria, VA
- Laskey, K. B. and Mahoney, S. M., 2000, "Network Engineering for Agile Belief Network Models", In IEEE Transactions on Knowledge and Data Engineering. Vol. 12, No. 7. July/August.
- Laskey, K. B., 1997, "Recognition Fusion: Network Fragments Inferring Entity Existence and Type from Sensor Reports", Report to IET.
- Li, Z., and D'Ambrosio, B., 1994, "Efficient inference in Bayes Networks as a combinatorial optimization problem", International Journal of Approximate Reasoning, 11:55-81, 1994.
- Longley, Goodchild, Maguire, and Rhind, *Geographical Information Systems, Volume 1, Principles and Technical Issues*, John Wiley and Sons, 1999.
- Mahoney, S., and Laskey, K. Constructing Situation Specific Belief Networks. UAI 1998: 370-377.
- Neopolitan, R.E., 1990, *Probabilistic Reasoning in Expert Systems*, New York, John Wiley & Sons, Inc.
- Pearson, A.R., and Wright, J.S., 1980, "Synthesis Guide for Cross-Country Movement", Engineer Topographic Laboratories Report No. ETL-0220, Ft. Belvoir VA.
- SEDRIS, 2000, "Synthetic Environment Data Representation and Interchange Specification (SEDRIS)", web page <http://www.sedris.org/>
- TEC, 1999, FOUNDATION FEATURE DATA PROTOTYPE EVALUATION, VOLUME 1 - EVALUATION, OVERVIEW AND KEY FINDINGS, briefing prepared by US Army Corps of Engineers, Engineer Research and Development Center, Topographic Engineering Center, Ft. Belvoir VA.
- Whittaker, J., 1990, *Graphical Models in Applied Multivariate Statistics*, John Wiley & Sons, New York
- Wright and Riggs, 1998, "Scene Generation Quality Assessment", SBIR Topic A97-165, Phase I Final Technical Report, produced for USATEC, 1998.
- Wright, E., (2002) "Understanding and Managing Uncertainty in Geospatial Data for Tactical Decision Aids", Ph.D. Dissertation, George Mason University, August 2002.

- Wright, 2000a, "Extending Techniques From "Standard" Error Theory To Categorical Data And Local GIS Operations", proceedings of the 2000 American Society of Photogrammetry and Remote Sensing (ASPRS) Conference Spring 2000, Washington, DC, May 2000.
- Wright, 2000b, "Probabilistic Models In Geographic Information Systems – Bayesian Networks For Management Of Uncertainty", proceedings of the 4th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences, Amsterdam, July 2000.
- Wright, 1999a, RTV Data Quality Implementation Plan, unpublished paper produced for the Rapid Terrain Visualization (RTV) Advanced Concept Technology Demonstration (ACTD) of the Joint Precision Strike Demonstration Project Office.
- Wright, 1999b, "Visualization of Uncertainty in GIS Data or Products", paper presented at the 1999 Joint Statistical Meeting, Baltimore MD, Aug 1999.
- Wright, 1998, "Application of Bayesian Networks for Representing Uncertainty in Geospatial Data", presented paper, American Society of Photogrammetry and Remote Sensing (ASPRS) Conference Spring 1998, Tampa, FL, Apr 1998.
- Wright, E. J., 1998, "Virtual Reality, or is it? Modeling Uncertainty in Virtual Environments", proceedings I/ITSEC 1998, Orlando FL, Dec 1998.
- Wright, 1997, "Information Integration for Remote Sensing", International Symposium on Spectral Sensing Research (ISSSR), San Diego, CA, Dec 1997.

Appendix A. Acronyms & Expansions

<i>Acronym</i>	<i>Expansion</i>
AI	Artificial Intelligence
CCM	Cross Country Mobility
COA	Course of Action
DARPA	Defense Advanced Research Projects Agency
DEM	Digital Elevation Model
DFAD	Digital Feature Analysis Data
DTED	Digital Terrain Elevation Data
DTSS	Digital Topographic Support System
EDCS	Environmental Data Coding System
FFD	Foundation Feature Data
GIS	Geographic Information System
IET	Information, Extraction and Transport, Inc.
IPB	Intelligence Preparation of the Battlefield
ITD	Interim Terrain Data
KAW	Knowledge Acquisition Workshop
KBMC	Knowledge-Based Model Construction
KBS	Knowledge Based System
KE	Knowledge Engineer
KR	Knowledge Representation
LOS	Line Of Sight
M&S	Modeling and Simulation
METT-T	Mission, Enemy, Terrain, Troops, Time available
MPRS	Mission Planning and Rehearsal System
NGA	National Geospatial-intelligence Agency
NIMA	National Imagery and Mapping Agency
OCOKA	Observation & fields of fire, Cover & concealment, Obstacles, Key terrain, Avenues of approach
PI	Principal Investigator

RKF	Rapid Knowledge Formation (DARPA program)
SME	Subject Matter Expert
TDA	Tactical Decision Aid
TEC	Topographic Engineering Center
TM	Theater Missile
TSKB	Terrain Suitability Knowledge Base
TTP	Tactics, Techniques, & Procedures
VITD	VPF format Interim Terrain data
VPF	Vector Product Format
WES	Waterways Experiment Station

Appendix B. User Recognition of the Need for Quality / Uncertainty Assessments

IET representatives attended three technical topographic / terrain related meetings sponsored by the US Army Topographic Engineering Center (TEC). Issues raised at these conferences revealed the growing user recognition of the need to include data quality / uncertainty in terrain analysis.

1. Operation Enduring Freedom (OEF) / Operation Iraqi Freedom (OIF) Lessons Learned Conference at US Army TEC, 23 Sept 2003. This meeting involved representatives from Army topographic units who presented briefings on activities during the Iraq conflict, including successes, and issues.

One issue identified and illustrated with several “war stories” was the problem of “pseudo” terrain experts. This problem occurs because many of the computerized command and control systems include powerful software for manipulating and displaying terrain data (for example FalconView and Arc View). The software makes it easy for any user to produce terrain products. Unfortunately these users lack understanding of the data (and its limitations), especially the data accuracies, and users would often reach unwarranted conclusions - which would disagree with the assessments produced by others, and especially by the terrain analysts in the terrain teams. One Terrain Warrant Officer said that he spent a large portion of his time as a “fireman” putting out “fires” created by these pseudo terrain experts. This problem is exacerbated by the state of terrain data availability which is very heterogeneous (data comes from a variety of sources, with a wide range of currency, coverage, resolution, quality, etc.)

- As a possible solution, there was a call for more training of potential users. But there are challenges with that solution: 1) potential users means almost anyone who has access to almost any computerized system. Training everyone would be very expensive. 2) Also, some of the biggest offenders were Engineer Officers, who have had more training than most (but apparently still not enough).
- Another solution option discussed is to somehow prohibit or prevent anyone who lacks appropriate training (that is, anyone who isn't a terrain analyst) from using these terrain applications. This is actually the preferred solution by most of the terrain people. But it's too late - the applications have already been fielded. The trend is for more and more of these terrain applications to be available on individual platforms and in many cases embedded (hidden) inside automated mission planning and command and control systems.
- A third solution option is to make the terrain application software smarter, so the software provides tools to make the user aware of the appropriateness / usability of the data for a particular task. This is exactly the approach that IET's prototype terrain suitability knowledge base has taken.
- This issue was a recurring theme over all 3 meetings.

Another issue was “access to Subject Matter Experts (SME)”. Even the terrain analysts are not expert in everything. And the training of terrain analysts is much less technical than it was even 10 years ago. One approach that worked well was “tele-engineering”. This is sponsored by the

Corps of Engineers, and makes civilian engineers and scientists at the Waterways Experiment Center, TEC, or other places, available to confer with deployed soldiers in the field via VTC. Tele-engineering has been used to provide expertise on hydrology, dams, bridges, soil dynamics, and other technical engineering topics. This has worked well in Bosnia, and also in Iraq. However, the experts are not available 24/7, nor are they available to everyone who needs them.

- This issue illustrates the importance of RKF technology to be able to provide soldiers in the field with access to Knowledge Bases.

2. Line of Sight (LOS) Technical Working Group Meeting, sponsored by TEC, 23 Sept 2003. The LOS Working Group is a recurring technical meeting that reports on and discusses application and algorithms for LOS.

One new application discussed is the need for, and potential algorithms to achieve, a Probabilistic LOS algorithm that takes into account information about elevation accuracies, as well as information about vegetation including estimates of understory, to generate a statistical estimate of LOS based on available data and the data's quality.

- This issues demonstrates the relevance of the probabilistic LOS algorithms / applications being developed for IET's prototype terrain suitability knowledge base.

3. Engineer Research and Development Center (ERDC) / TEC Technical Exchange Meeting, 24/25 Sept 2003. This is a recurring meeting that provides a forum for technical presentations on ongoing research and developments in the topographic community.

Dr. Paul Krause, TEC, presented interim results of a study to develop prediction equations to predict parameters of vegetation understory from vegetation overstory data. For example, understory is very important to terrain applications like Cross Country Mobility (CCM), LOS, and weapons engagement ranges, but understory data is typically not available in standard databases. (In large part this is because it is too difficult and expensive to collect). Overstory data is generally available. The most important overstory parameter is vegetation height. The application of prediction equations would provide predicted understory data to support sophisticated CCM and LOS algorithms, as well as advanced simulations / visualizations. Dr. Krause has done some limited field data collection, and has some regression equations that fit his data fairly well ($r^2 = .8$), although he envisions his results being used as deterministic equations. He now has funding to do more data collection designed to collect data from 15 of the 21 world Biomes. A report describing this work is available.