A Knowledge-Based Tactical Decision Support System Integrating Terrain and Weather Data

L. A. Staub, M. J. Carlotto, B. G. Lee and R. A. Upton*

The Analytic Sciences Corporation
One Jacob Way
Reading, MA 01867

A distributed, knowledge-based system which integrates terrain, weather, threat and military organization/doctrinal information for real-time decision support in a tactical battlefield environment is described. The system consists of data fusion, terrain and weather analysis modules, together with a rule-based system for inferring trafficability and cross-country mobility. Within the data fusion module, multi-spectral, black-and-white and weather satellite imagery, digital terrain elevation and collateral data are registered to a common reference base. The terrain and weather analysis modules then generate a physical description of current terrain and weather conditions from the fused source imagery. Using this data, trafficability and cross-country mobility can be inferred for selected classes of wheeled vehicles.

Key Words: Multi-source data fusion, rule-based image interpretation systems, precipitation estimation, surface material classification, terrain analysis, image registration, decision aids, decision support tools.

1. INTRODUCTION

The time in which a tactical battlefield commander has to make significant decisions concerning the deployment of his forces, allocation of reconnaissance and targeting resources, and responding to threats is decreasing rapidly. This is due in part to the increased capabilities of processing, communication and sensor technologies to provide him with substantial quantities of information in real-time. As a consequence, automated techniques for filtering and fusing the variety of data presented to him, and for assisting him in making critical decisions are needed.

This paper describes a prototypical knowledge-based system, termed the Very Intelligent Exploitation Workstation (VIEWS), for providing such tactical decision support. VIEWS integrates terrain, weather, threat and military organization/doctrinal information to support dynamic, real-time decisions centering about cross-country mobility, line-of-communication, trafficability, observability, avenues of approach, possible staging areas and target locations, and optimal search strategies. VIEWS currently focuses on providing a dynamic description of terrain, in particular, estimating the cross-country mobility of military vehicles as a function of terrain composition, slope and weather. Cross-country mobility (the degree to which a vehicle can move and maneuver away from all-weather routes) is an important element in what is known as the intelligence preparation of the battlefield (IPB) which attempts to provide a graphic presentation of the battlefield and possible enemy courses of action. The IPB process [5,6] is currently very labor intensive and requires the direct participation of terrain, weather and other experts. As a consequence, it is a good candidate for the application of expert systems technology. There exist, however, few published reports on the application of artificial intelligence, (in particular, expert systems) to IPB [2,3].

The objective of the present VIEWS effort is to develop an IPB decision support system capable of exploiting currently available imagery and collateral databases in order to address the problem of determining cross-country mobility for selected classes of military vehicles. Our approach is based on the application of image processing/understanding techniques to build a physical description (i.e., terrain and weather) of a selected area of interest using imagery and collateral data sources. Knowledge-based techniques are subsequently applied to infer tactically-significant information as a function of weather, terrain and threat. The scope of this paper is limited to that part of the IPB process which relates to assessing cross-country mobility.

The organization of the remainder of the paper is as follows. Section 2 provides an overview of the current VIEWS system which is divided into modules which perform data fusion, and terrain, weather and cross-country mobility analysis. Our implementation of the cross-country mobility rule base and inference engine is then described in Section 3. An example illustrating the use and capabilities of VIEWS is provided in Section 4. Finally, future extensions to VIEWS are discussed in Section 5.

*Currently at PSC 10340 Democracy Lane, Fairfax, VA 22030

2. SYSTEM OVERVIEW

The decision support system is divided into four modules: data fusion, terrain analysis, weather analysis, and cross-country mobility analysis as shown in Figure 1. The data fusion module registers real-time imagery, external databases and collateral data sources to a reference map (e.g., a digitized topographic map), and segments the area of interest into regions having similar terrain and weather conditions.

The terrain analysis module supplies the system with static information on surface material composition (SMC), soil type, and slope. This information is provided in a map-like form and is derived from imagery and external databases. Knowledge of SMC allows land use/cover categories, hydrologic features and transportation routes to be identified, and soil type and conditions to be inferred when soil maps are not available. Soil type is an important factor in estimating soil moisture and soil trafficability under various weather conditions. Slope, in conjunction with the above information, is needed to infer the cross-country mobility of particular vehicle types.

The weather module analyzes past and present weather information to infer the current meteorological context for IPB analysis. Current weather information is provided in the form of precipitation maps derived from weather radar and meteorological satellite imagery and textual reports generated by weather stations. Weather impacts mobile operations in a variety of ways. For example, the trafficability of terrain is reduced by increased soil moisture, accumulated snow cover and icing conditions. Alternately, bodies of water and wetland areas become trafficable as temperatures drop and the freezing period lengthens. Transit operations are hampered by low visibility resulting from fog and heavy precipitation.

Analysis of cross-country mobility (CCM) for vehicles is performed by the CCM module which interprets terrain and weather conditions and vehicle characteristics in order to assign a cross-country mobility rating to an area (e.g., go/no-go, or good/moderate/severe). The CCM assessment for a geographic area is primarily a heuristic process characterized by the need to identify the key factors which affect CCM. In addition, the knowledge required to interpret the impact of weather and terrain on CCM is diffuse and more easily represented in the form of situation-action rules than algorithmically. In this context, a situation describes the key weather and terrain conditions which are required to trigger a CCM assessment. The heuristic nature of the knowledge and the need to organize the diverse weather and terrain knowledge into manageable units favors a rule-based implementation.

3. CROSS-COUNTRY MOBILITY RULE-BASED SYSTEM IMPLEMENTATION

This section discusses the design and implementation of the cross-country mobility rule-based component of the VIEWS system. The architecture shown in Figure 2 contains a segmentor which partitions the area of interest into regions with similiar terrain and weather conditions, display functions for viewing CCM and intermediate map representations, and a rule-based system. The latter is divided into rule components, each of which contains a ruleset, a context in which rules are applied and a control mechanism which applies rules within a context. The initial implementation described in this paper contains four rule components which interpret weather, terrain, and vehicle characteristics for assessing soil moisture, soil trafficability, slope and cross country mobility.

3.1. Composite Region Representation

Once static terrain (slope percentage, surface material composition and soil type) and current weather conditions (temperature and precipitation) have been determined, the geographic area is segmented into regions having similar characteristics; e.g.,

SMC = forest
10% < slope percentage < 20%
soil type = clay
precipitation = 0
temperature < 32^{o} F.

For each kind of region in the image an instance of a *composite region token* is created. (Tokens represent significant areas in images and are implemented using flavors in Zetalisp.) Tokens contain slots for the above terrain and weather conditions, for information inferred by the rule-based system, and for ancillary information such as what rules have fired during prior evaluation. Transforming registered slope-percentage, surface material, and soil maps (each 540x1024 pixels in area) into composite regions typically produces a list of 300-400 tokens, effectively reducing the dimensionality of the data by several orders of magnitude.

3.2. Knowledge Representation

Domain knowledge about the effects of terrain and weather on the mobility of vehicles is represented as production rules with the typical structure of left hand side (LHS) conditions and right hand side (RHS) conclusions. Conditions are stated as the logical conjunction and/or disjunction of clauses which are predicate-attribute-value tuples. Right hand side conclusions specify assertions to be made in the current context or processes to be executed.

The knowledge used to provide cross country mobility assessments is distributed both horizontally and vertically within the system. Knowledge about the relevance, utility, and inferencing strategies of a rule base is vertically distributed to higher levels of control within the system (Figure 3). Currently, this vertical structuring of knowledge supports the implementation of both general and ruleset-specific strategies which control rule execution.

Rules are distributed horizontally across rulesets based upon their conclusions. For example, rules which make statements about soil moisture make up one ruleset, while rules which draw conclusions about cross-country mobility are segregated into a separate ruleset. Distribution of rules into separate rule bases is advantageous since it facilitates incremental knowledge engineering of rules, allows useful generalizations to be made about a ruleset, and provides greater efficiency during the match cycle.

Distributing domain knowledge across rulesets supports a top-down and incremental approach to knowledge engineering. For example, the cross-country mobility ruleset (the top-level interpretation in the current implementation) was developed first, focusing on the kind of information and assessments needed for inferring cross-country mobility. It was only after an initial version of this ruleset was completed that attention was focused on developing the necessary intermediate interpretations implied by the cross-country rule conditions (LHS clauses). In this way, soil trafficability, soil moisture, and slope rulesets were developed. Experience with the current implementation of the system has demonstrated that following this top-down, incremental knowledge engineering strategy facilitates the rule writing and debugging process. The vocabulary of a ruleset is kept at a relatively high level (thus rules are easier to write) and the number of interacting clauses within a rule is minimized (thus rules are easier to debug) as shown in Figure 4.

Since a ruleset is organized around a single goal (i.e., to make a certain type of assessment or assertion), generalizations about the ruleset can be made and used elsewhere within the system. This allowed knowledge about a ruleset to be explicitly and centrally represented and used at a higher level of interpretation, for example, in the specification of ruleset-specific conflict resolution strategies (Section 3.4).

Another advantage of the distributed rule components is the greater efficiency gained in the rule matching process. This efficiency results from having fewer and less complicated rules within a ruleset and from having the ruleset organized around a common goal. Rules not directly relevant to the satisfaction of the current goal are not considered within the current recognize-act cycle.

3.3. Rule Context

The *context* in which a rule is tested and fired has a frame-like representation and a function analogous to short-term or working memory [1,7]. A context provides a locus and structure for consolidating the data required for the rule based interpretations. Information within a context is represented as slots and their associated value(s) or attached procedures.

The creation of a context is triggered by the application of a ruleset. The data which is organized into a context comes from several sources, including a composite region instance, global and default data, and other frame-like constructs which contain information about vehicle characteristics and military organization and doctrine. The attributes and values which appear in a rule's LHS clauses are the same ones which appear as slots and values within a context. The effect of firing a rule (i.e., executing the RHS clauses) is to modify some slot(s) within the current context. These conclusions or inferences are communicated to other parts of the system via composite regions.

3.4. Control

Execution of rules is controlled using a mixed strategy including both data-driven and goal-driven control. The match and execution of rules within a ruleset (i.e., the recognize-act cycle) is a data-driven process. At a more global level however, the invocation of an entire ruleset is controlled using a goal-driven strategy. For example, if the goal is to compute cross-country mobility of vehicles, the cross-country mobility ruleset is activated. Prior to actually entering the recognize-act cycle for this ruleset, the ruleset is queried for the categories of information needed for its inferencing (e.g., surface material class, accumulated precipitation, soil trafficability, slope). This information is collected either from existing sources (imagery, data bases, global declarative knowledge) or it is computed by another rule-based component. For example, the activation of the cross-country mobility ruleset leads to the invocation and application of other rulesets such as slope and soil moisture. It is only after all data collection and relevant inferences have been made (or been excluded in some cases as unavailable or too costly to collect), does forward chaining within a ruleset (in this example, the cross-country mobility ruleset) take control.

A further element of control exists within the recognize-act cycle at the point at which a number of rules have been activated (i.e. their LHS conditions have matched within the current context) and a decision must be made about which rule(s) to fire (conflict resolution). Meta-level knowledge about the ruleset's domain is invoked at this time. For example, knowledge about the weather and terrain factors used by cross-country mobility rules is employed in the conflict resolution strategy for deciding which rule(s) within the set of active rules has the greatest relevance or the highest priority in the current context. Knowledge about the relative impacts of steep slopes, dense forest cover, soil trafficability and weather conditions can be used to favor rules which infer cross-country mobility based on the steepness of the slope over rules which infer cross-country mobility based on soil trafficability since the relative trafficability of the soil is not relevant when there exists a steep slope. Likewise, rules which identify blizzard conditions may be favored over rules which mention terrain factors since zero visibility may be seen as the primary restriction to cross-country mobility.

If, after applying a domain or ruleset specific strategy, there still remain several active rules, a more general strategy is invoked. One of the strategies employed for all rulesets favors the rule which uses the greatest number (diversity) of weather and terrain factors. This is based on the heuristic that more information will generally lead to a better evaluation of a situation than less information. Other strategies may favor rules whose premises are based on (matched with) data of greatest reliability and accuracy or rules which, through experience, have proven to have a high degree of utility.

4. EXAMPLE

A 5.4 x 13.2 km site in Putnam County, NY was selected for developing the terrain and weather analysis modules and cross-country mobility related rulesets. This particular site was chosen based on its topography and the availability of imagery, terrain databases and collateral sources of sufficient quality to support initial experiments. Among the source materials used were:

- USGS 1:24,000 scale topographic maps
- USGS digital elevation models (DEMs)
- Deadalus AADS 1268 airborne multispectral scanner imagery
- SPOT-simulation imagery
- Soil Conservation Survey (SCS) soil maps
- Geostationary Operational Environmental Satellite (GOES) imagery
- Ground-based weather radar (New York City).

Systeme Probatoire d'Observation de la Terre is a multispectral sensor to be launched by the French late in 1985. SPOT-simulation imagery is a simulated version of the anticipated satellite product produced from Daededus AADS 1268 aircraft scanner data by SPOT-image Corporation. SCS soil maps were digitized from hardcopy. GOES imagery and NYC weather radar were obtained from the National Oceanographic and Atmospheric Administration (NOAA).

4.1. Image Registration

For precipitation estimation, GOES east and NYC ground weather radar were registered to World Data Base II using a polynomial warping technique. A similar technique was used for registering SPOT multispectral imagery, the SCS soil map and a slope map computed from the DEM. Ground resolutions ranged from 10 meters for sharpened SPOT-simulation imagery, to approximately 85x47 meters for the DEM, to 1 km (visible) and 8 km (IR) for the GOES imagery. Lower resolution imagery was interpolated to 10 meters.

4.2. Precipitation Estimation

A statistical classification approach was used for estimating precipitation rate from GOES imagery [8]. Weather radar from a location within the satellite coverage was registered to the GOES image and the different precipitation levels in the radar were used to indicate the corresponding rain areas in the clouds. Samples of no rain, light rain (.01-.65 in/hr) and heavy rain (0.66-1.5 in/hr) areas were extracted from one image. The statistics from these samples were then used to classify pixels into the three rainfall classes throughout the sequence of images for that day. Once instantaneous precipitation is computed for each image, the accumulated precipitation is calculated by summing the rainfall over a period of time.

4.3. Terrain Analysis

Statistical techniques were used for surface material classification. A statistical classifier, similar to the one used for precipitation estimation, was used for surface material classification. Six major SMCs were extracted in the SPOT-simulation image: water, forest, sparse vegetation, bare soil, crops, and concrete. In addition to SMC, soil type and slope percentage are required for inferring cross-country mobility. Soil characteristics were acquired from SCS maps and reports for Putnam county. Slope percentage (100 x rise/run) is computed directly from the DEM by spatial differencing. Relative slope (steep/not steep) is determined as a function of slope percentage, surface material composition, and current weather conditions.

4.4. User Interface to Cross Country Mobility Analysis

A major consideration in developing the CCM module was to provide the IPB analyst with the capability to interactively control and examine the CCM analysis process. Interactive control is provided by allowing the user to modify weather parameters (the most dynamic aspect of the problem) and apply individual rulesets to the composite region map. By modifying the weather parameters and subsequently applying a rule base within the context of these modifications the user can explore the impact of different weather conditions upon specific interpretations like soil moisture, soil trafficability, or cross country mobility. The user is further able to examine the CCM interpretation process by selecting specific geographic regions (displayed as a black and white aerial photo or a CCM rating map on a color monitor) and having all information known about this region displayed at the terminal.

Another major consideration in the development of the system was the IPB requirement for graphic presentation of all phases of the fusion and analysis process. Within the design of decision support aids this requirement motivated the inclusion of graphic representation and presentation techniques for top-level analyses such as CCM, intermediate interpretations such as soil moisture, and basic data inputs such as soils, elevation, surface material classes and black-white imagery.

The graphic representation of results impacts the knowledge-based component of the system in two important ways. First, the conclusions or inferences derived from rule-based analysis must have a representation which can be translated by display functions to graphic presentation on a color monitor. This was accomplished by including slots for these results within the composite region token which is used by all the display functions. Secondly, it was necessary to keep an explicit representation of the intermediate inferences (soil moisture, slope, soil trafficability) which are used to support cross-country mobility analysis. This was accomplished by having each rule component post it conclusions within the composite region token in the same way as the CCM results are included (Figure 5).

Consolidation and registration of all data and interpretations in the composite region map provided a uniform mechanism for the display of all phases of the CCM process. Associating an audit trail or trace of the inferencing process with this representation further provided a method for graphically examining and tracking the analysis process from end interpretation to initial data inputs for each and any region. This graphic presentation of information and interactive cabality to select regions for examination is useful both as a decision support tool for the IPB analyst and as a tool for the knowledge engineer during rule development.

5. SUMMARY

A prototypical system for generating decision aids from current imagery and databases for IPB has been described. The emphasis has been on developing techniques for multisource fusion and representation, for building physical descriptions of an area of interest in terms of its terrain and weather, and for inferring CCM using a knowledge-based approach. Future work involves refining the current rulesets, further automating terrain and weather analysis (e.g., using knowledge-based techniques such as those used in MSIAS [4]), and adding spatial reasoning components for identifying lines of communication and avenues of approach.

References

[1] A. Barr and E. A. Feigenbaum, *The Handbook of Artificial Intelligence (Vol. 1)*, William Kaufman, Inc., Los Altos, CA, 1981, pp 190-199.

[2] R. P. Bonassa, "Expert systems for intelligence fusion," *Proceedings of the Army Conference on Application of Artificial Intelligence to Battlefield Information Management*, White Oak, MD, April, 1983, pp 101-116.

[3] T. Campen and D. E. Gordon, "Application of artificial intelligence to tactical operations," *Proceedings of the Army Conference on Application of Artificial Intelligence to Battlefield Information Management*, White Oak, MD, April, 1983, pp 81-100.

[4] R. D. Ferrante, M. J. Carlotto, J. M. Pomarede and P. W. Baim, "Multi-spectral image analysis system," *The First Conference on Artificial Intelligence Applications*, Dec 5-7, 1984, Denver, Colorado.

[5] IPB, Intelligence Preparation of the Battlefield, SUPR 66000-A, June, 1983, U.S. Army Intelligence Center and School, Ft. Huachuca, AZ (uncl).

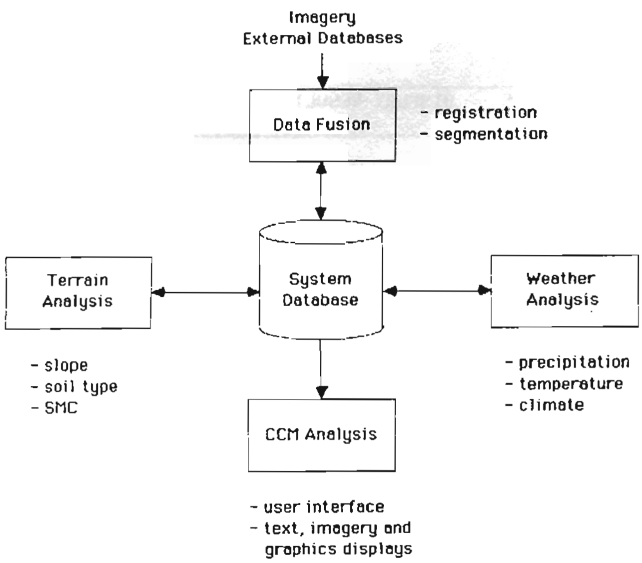

Fig. 1 VIEWS Functional Organization

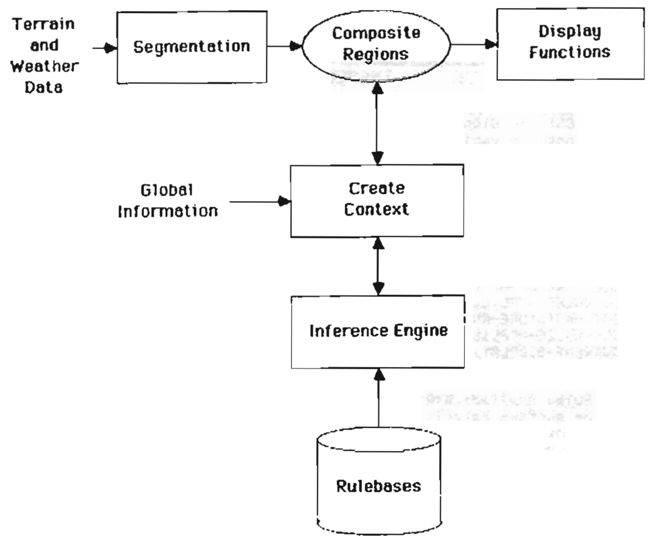

Fig. 2 Architecture of CCM Rule-Based System

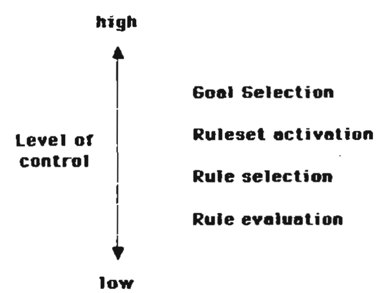

Fig. 3 Vertical Distribution of Control Knowledge

Cross Country Mobility Rule

IF surface material is bare soil and
soil moisture is wet and
soil trafficability is moderate or poor
THEN
cross country mobility is severe.

Soil Trafficability Rule

IF soil content is clay or silt and
soil drainage is very poorly drained and
soil bearing strength is moderate
THEN
soil trafficability is poor.

Soil Moisture Rule

IF accumulated precipitation is heavy rain and
current precipitation is light or heavy rain
THEN
soil moisture is wet.

(a) Distributed into three rulesets

IF surface material is bare soil and
accumulated precipitation is heavy rain and
current precipitation is light or heavy rain and
soil content is clay or silt and
soil drainage is very poorly drained and
soil bearing strength is moderate
THEN
cross country mobility is severe.

(b) Consolidated into one ruleset

Fig. 4 Distribution vs. Consolidation of Knowledge

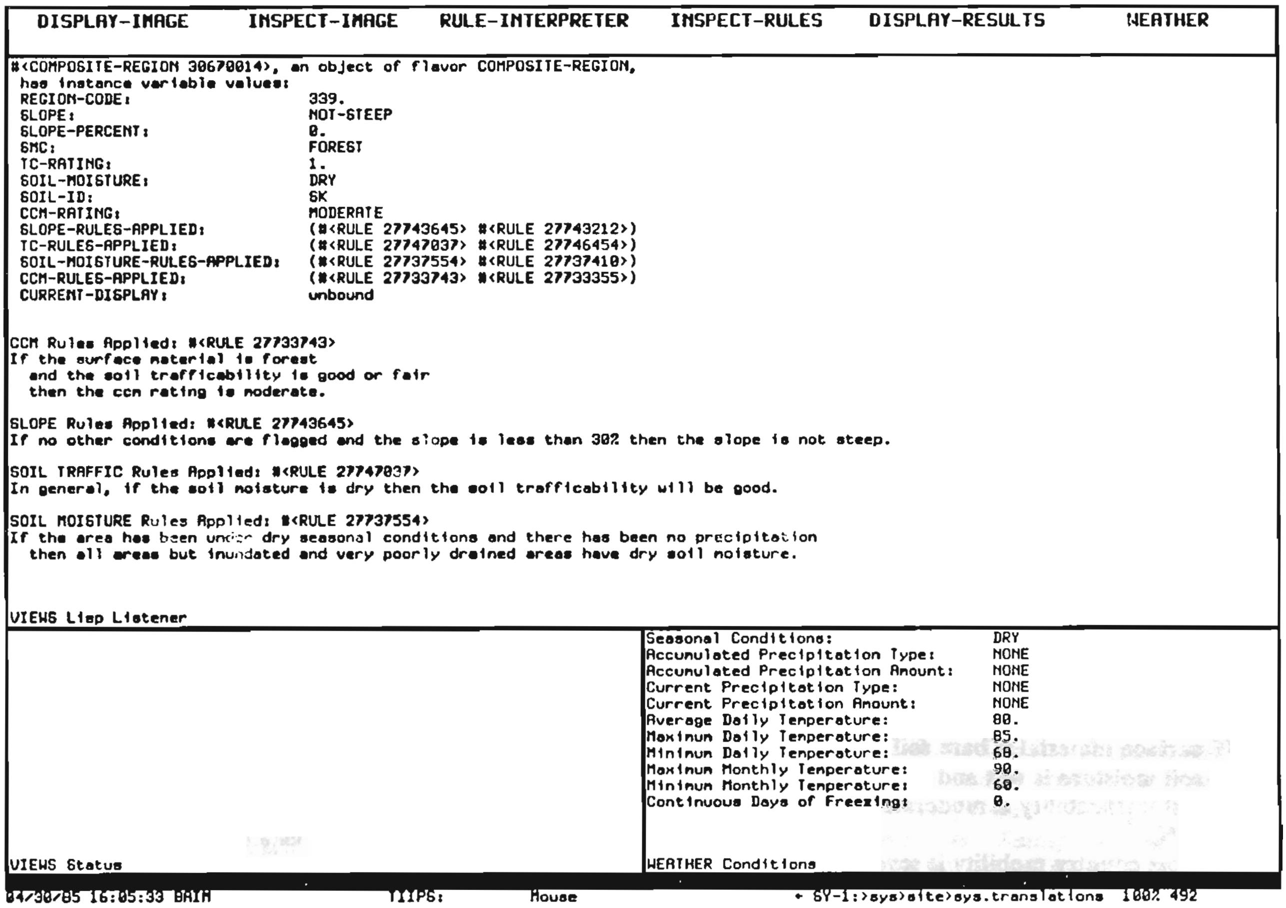

Fig. 5 VIEWS User Interface