66

VOLUME 1, NUMBER 3, 1986

HEALTH CARE INSTRUMENTATION

THE INTERNATIONAL JOURNAL OF CURRENT MEDICAL TECHNOLOGY

CONTENTS

EXCERPTA MEDICA / An Elsevier Company

COMPUTERIZED NEUROLOGICAL DIAGNOSIS: A PARADIGM OF MODELING AND REASONING

ZHIGANG XIANG, JERRY G. CHUTKOW,† STUART C. SHAPIRO,* and SARGUR N. SRIHARI**

*Department of Computer Science.and †Department of Neurology, School of Medicine, State University of New York at Buffalo, Buffalo, New York

SUMMARY

Neurological diagnosis is an ideal application of artificial intelligence. This paper reviews work on the design of model-based expert computer systems intended to simulate human behavior, with an emphasis on the modeling schemes and reasoning mechanisms of these systems, and presents the authors' experimental approach to neuroanatomic localization.

INTRODUCTION

Computer-aided diagnosis, an active area of research in computer science, has potential applications in domains ranging from machine repair to medicine. A diagnostic expert system (DES) is a computer program capable of inferring in an expert-like manner the internal status of the subject of a diagnostic analysis from the subject's observed behavior and supplementary data. The design of DESs, a branch of artificial intelligence, incorporates such techniques as symbolic representation, pattern-directed inference, and heuristic reasoning.[1]

How the knowledge involved in diagnostic decision-making is conceptualized determines the structure and function of a DES. For example, one can map behavioral changes (symptoms, physical findings, adjunctive laboratory abnormalities, or any combination of the three) to symbolic names for the underlying cause of the internal problem (eg, specific diseases or pathophysiological states). Mapping is accomplished by a combination of categorical and probabilistic reasoning that proceeds through intermediate stages to reach diagnostic conclusions. This approach is used in many systems, including INTERNIST-1, PIP, CASNET, MYCIN, and PUFF.[2] Such systems are based on "shallow" knowledge, because they do not relate input data and conclusions to the structure of the subject, to the function associated with components of structure, or to other knowledge such as taxonomy.

A system that associates behavioral changes with a functional model of the subject has "deep" knowledge. When this core representation is coupled with reasoning strategies,[3,4] the system is based on "first principles."[5] Understanding a subject's structure and function allows a DES to explain the causal relationships between internal status and behavior. For instance, a machine repair system needs to know how the machine works in order to analyze complex malfunctions, whereas a knowledge of how the machine fails makes it possible to handle simple disorders by mapping symptoms to specific types of dysfunction.

Efficient and economic diagnostic reasoning depends on the appropriate selection of models. A subject can be modeled logically, physically, or in different ways at various levels of abstraction, depending on the required diagnostic tasks. Logical modeling provides a function-oriented abstraction of the subject in which each logical component tends to have its localized contribution to the overall functional behavior of the subject. The knowledge base of CADUCEUS,[6] which exemplifies the use of logical modeling, includes causal relations between observations and diseases (also pathophysiologic states), nosology, and heuristic information.

Physical modeling provides a structure-oriented abstraction in which each physical component tends to have its concrete appearance as a part of the subject. Two models of the same type have a hierarchical relation if one is an abstraction of the other.

The logical and physical models of a subject usually correspond. Thus, generally speaking, the power train of an automobile consists of two functional (logical) components: the engine, providing power, and the transmission, converting the power to motion. These, in turn, correspond to the two physical components, the engine assembly and the transmission assembly. The diagnosis of faults in digital circuits[4,5] also relies explicitly on logical modeling (the circuit diagram) as well as physical modeling (the chip arrangement, which accounts, for example, for heat or magnetic field disturbance). The two types of models sometimes intersect. A digital circuit may be functionally modeled by a logical

0883-0762/86/$3.50

diagram in which such logical components as "AND gates" or "OR gates" contribute locally to the whole circuit, but several logical components that are functionally irrelevant may be enclosed in the same chip (ie, correspond to the same physical component).

The physical structure of a subject refers to its organization in three-dimensional or, in special cases, two-dimensional space. The representation of three-dimensional structure has been studied for a long time in computer science.[7,8] The methods used are either analogical or propositional.[9]

Analogical representations are detailed geometric descriptions that allow objects to be specified succinctly and unambiguously by means of coordinates. This specification can be accomplished by mathematical equations or by dividing three-dimensional space into volume elements (voxels), where sets of voxels specify the curves, surfaces, or objects within the space.[8] The advantages of these approaches are that inference rules for spatial reasoning can be implemented by algorithms from computational geometry and that graphics and image processing techniques can be adapted without much difficulty. The disadvantage of these approaches is that they usually overspecify the real world. Every entity cannot or need not be given an exact geometric description. Furthermore, these approaches sometimes have little relation to the cognitive approaches of human beings.

Propositional representations abstract salient topological features in order to describe entities in terms such as shape and position and to describe spatial relationships by adjacency, connectivity, direction, and other terms. Propositional representations, which favor modeling intelligence, may overcome some of the disadvantages inherent in analogical representations. They, too, however, have limitations: not all structural information can be expressed in the form of propositions. Sometimes, spatial information is better depicted in pictures.

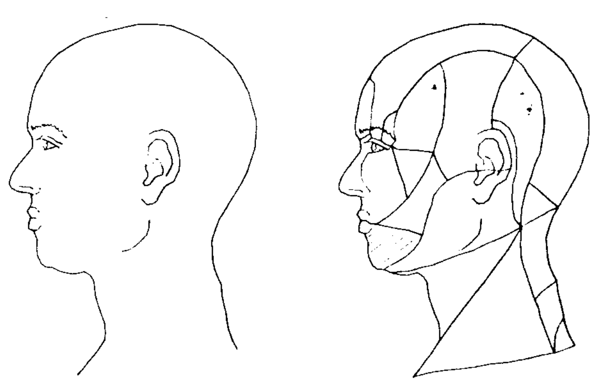

Figure 1. Head and neck drawn on a graphics screen *(left)*, with superimposed distribution of cutaneous sensory nerves to the affected area *(right)*. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government*. Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

In neurological diagnosis, reasoning processes rely heavily on both physical and logical modeling. Even early diagnostic programs, which do not use techniques of artificial intelligence, maintain simple models of the neuroanatomy.

The following sections of this article describe the fundamentals of neurological diagnosis, review and comment on major publications on computerized neurological diagnosis, with an emphasis on modeling schemes and reasoning mechanisms, and describe NEUREX, our experimental approach to the integration of modeling and diagnostic reasoning.

FUNDAMENTALS OF NEUROLOGICAL DIAGNOSIS

Neurological diagnosis[10,11] is the process by which diseases of the neurologic system (brain, spinal cord, peripheral nerves, neuroeffectors and neuroreceptors, and supporting structures) are identified. It is ongoing serial acquisition and analysis of data about a patient (subject).

In the first stage of the clinical database, the neurologist collects preliminary data, including qualitative and quantitative descriptions of symptoms (Sxs), the relationships between Sxs, the results of past and present physical examinations (Pxs), adjunctive laboratory data (Lxs), other relevant information, and the overall temporal profile or course of the illness. The clinician documents these (particularly Pxs) in writing on forms and on pictorial drawings. The latter not only indicate the extent of the disability but, when designed appropriately, also provide considerable information about the anatomy underlying the findings. For example, by indicating sensory losses as in Figure 1 (right), the neurologist can readily identify malfunctioning nerves.

In the next stage—neuroanatomic localization—the diagnostician uses functional general anatomic and functional neuroanatomic knowledge to infer the presence and site(s) of the cause of the neurologic Sxs, Pxs, and Lxs. This stage, providing the scientific foundation for the remaining diagnostic analysis, consists of two steps: axial and transverse localization. In the former, clinical data are assigned to their appropriate axial neurosystem(s).

"Axial neurosystem" is an idiosyncratic phrase conceptualizing a neuroanatomic-physiologic or neuroanatomic-psychologic unit transmitting or processing a specific set of clinically definable functions and roughly paralleling the axial lines of the body and limbs as it extends physically through many of the transverse segments into which the neurologic system is divided. A neurosystem or one of its individual components may transmit impulses related to specific sensations; integrate and evaluate sensory data; store information; regulate consciousness, attention, or affect; program motor, communicative, or secretory responses; or effect the desired responses. A neurosystem, depending on its functional state, gives rise to different types of Sxs, Pxs, and/or Lxs. Therefore, these data (Sxs, Pxs, and Lxs) can be used to infer the status of the neurosystem.

An axial neurosystem, in turn, is subdivided transversely into two major segments: one part transversing the parenchymal central nervous system (brain and spinal cord) and the other part passing through the peripheral nervous system (the parenchymal neurologic system external to the brain stem and the spinal cord). The central nervous system (CNS) and peripheral nervous system (PNS) are further subdivided transversely, in a manner described below, thereby setting up precise spatial coordinates with which to pinpoint a lesion.

For neuroanatomic localization, the neurologist or neurosurgeon must have in mind a functional three-dimensional model of the clinically important parts of the neurologic system along with its receptors (including eyes, ears, and sensors in the skin) and effectors (including skeletal muscles and sweat glands). Figure 2, a schematic of the spinal cord at three of its 31 levels, shows the distribution of a few axial neurosystems (tracts) and segment-limited (transverse) structures. Half of the cord mirrors the other half, and the physical positions and, therefore, spatial relations are fixed. Although the physical positions and spatial relations vary from one individual to the next, the variations are relatively minor from a diagnostic point of view and can be handled statistically.

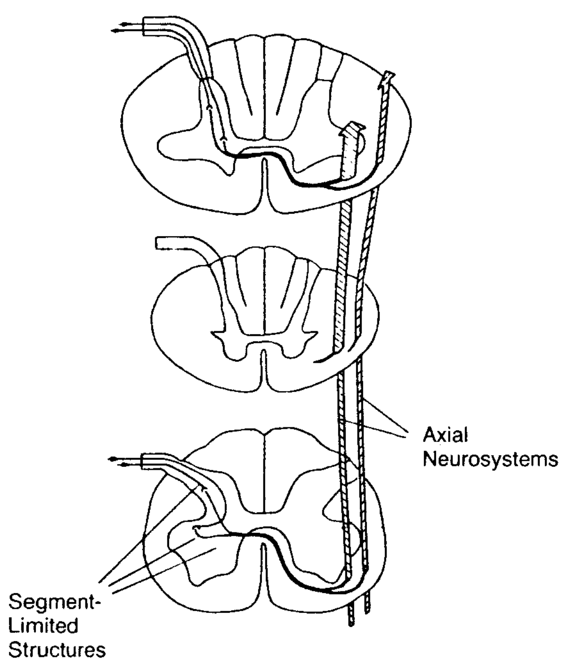

Figure 2. Schematic of the spinal cord at three levels: cervical *(top)*, thoracic *(center)*, and lumbar *(bottom)*. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

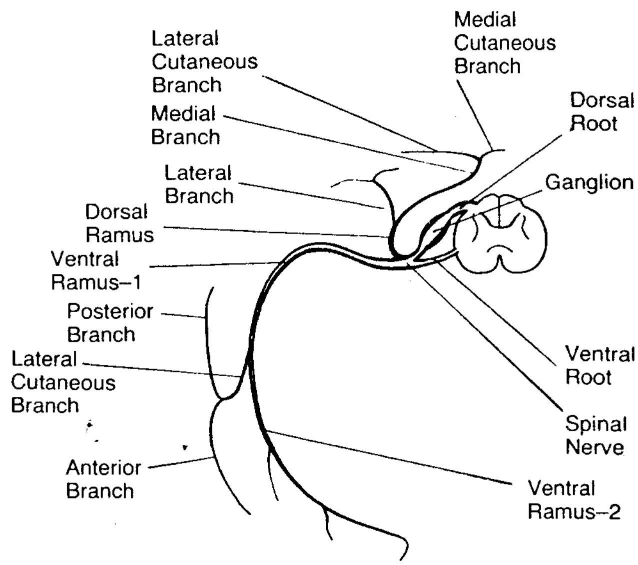

Figure 3. Schematic of a peripheral root system with the roots attached to a transverse segment of the spinal cord and a peripheral nerve system beginning with the spinal nerve. The two types of peripheral neurosystem distribution are the same unless they become part of a plexus. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

The axial neurosystems and segment-limited structures carry information about homologous areas on each side of the body. The homologous areas, too, are relatively consistent from one individual to another. For example, lesions on one side of the thoracic cord at a given level, regardless of etiology, cause a predictable pattern of neurologic deficits in different patients—provided that the lesions involve the same axial and segment-limited structures. Conversely, a combination of Sxs, Pxs, and Lxs can be traced back to an anatomic source, which helps the clinician decide whether the patient has a single well-circumscribed (focal) lesion, more than one focal lesion (a multifocal problem), dysfunction of one or more axial neurosystems (a systems-limited disorder), or an uncircumscribed process randomly involving many structures, usually in more than one transverse segment (a diffuse disorder). Similar anatomic principles govern structure, function, and localization in the PNS (Figure 3) and, with much greater complexity, in the brain.

Having determined the probable site and number of lesions, the clinician combines the neuroanatomic localization with elements of the clinical database (relationships between Sxs, Pxs, and Lxs; modifying factors; coexistence of other diseases; and the temporal profile of the illness) to deduce the underlying pathophysiological mechanisms (eg, ischemia or inflammation). Disordered anatomy and physiology are combined to form patterns suggesting pathogenetic categories of illness (eg, genetically determined disorders, vascular disorders, neoplasia). The pathogenetic categories of illness and multiple epidemiologic facts allow the clinician to concentrate on a small number of disease-specific etiologies, such as atherostenotic occlusion of a specific blood vessel, embolic infarction, or syphilitic endarteritis.

Hypothesis formation and data generation are ongoing, interactive, and mutually correcting. Hypotheses lead to a search for additional data and correction of erroneous information. These data and corrections, in turn, enhance the statistical probability of a particular anatomic localization, a particular pathogenetic category, or a specific disease while decreasing the probability of competing alternatives. Successful diagnostic analysis provides a basis for rational therapy.

PREVIOUS WORK ON COMPUTERIZED NEUROLOGICAL DIAGNOSIS

Previous work on computerized neurological diagnosis, ranging from simple programs to DESs, can be classified by the reasoning mechanism used to reach diagnostic conclusions and the modeling scheme (either analogical or propositional) used to represent neuroanatomy.

Approaches Based on Branching Logic

The most natural and straightforward way to utilize digital computers to aid diagnosis is to encode the relevant information by certain numerical data structures (eg, arrays) and the diagnostic decision-making processes by branching-logic algorithms. This approach is well supported by conventional programming languages, including FORTRAN, and by software designs including top-down programming and stepwise refinement (the process of dividing a complex task into several well-defined subtasks and solving each subtask in the same manner until primitive tasks, whose answers are known, are reached).

There are two basic methods with which to handle the "solution space" (the collection of all possible solutions). One is elimination of the solution space, and the other is evaluation of the solution space.

Elimination of the solution space consists of a step-by-step reduction in alternatives as critical ("clinching") data are collected. The diagnosis remaining at the end of the clinical inquiry is, by exclusion, the most likely one. Each step can be implemented by branching into the part of the program corresponding to the remaining solution space. Freemon[12] and Vastola and associates[13] have developed simple programs using this technique, without incorporating anatomic models, to classify patients according to their answers to a set of prearranged questions. Others have used the approach to guide the collection of clinical data in a systematic manner (eg, to produce case reports[14] and to record a simplified neurologic examination of patients in an intensive care unit[15]).

Evaluation of the solution space arrives at a diagnosis by accumulating supporting data for each possible answer. The cognitive significance of clinical data is quantified by arbitrary weights or points, the diagnosis with the greatest support at the end of the inquiry being the most likely one. This approach simulates the way people confirm or dismiss all possible choices by surveying every bit of available information.

Meyer and Weissman[16] have combined evaluation of the solution space with a rough analogical model of the brain stem, that portion of the CNS connecting the forebrain with the spinal cord, to localize lesions. This program maintains a representational array of ten transverse sections of the brain stem, each of which is divided by a 10×10 grid to produce a total of 1,000 volume units. Each unit has a malfunction factor, M, which indicates the number of involvements of the unit in malfunctioning pathways (analogous to axial neurosystems and functional structures limited to one transverse segment), and a function factor, F, which indicates the number of involvements of the unit in normally functioning pathways. Initially, each M and F is zero. The outcome of each clinical neurological test is used to increase the M and/or F of each unit involved. Finally, a modified net malfunction factor (equal to $M - F$ if $M - F > 0$ or 0 if $M - F \leq 0$) is calculated for each unit. The unit with the highest modified net malfunction factor is most likely to contain an anatomic lesion.

Francis and associates[17] adopted a similar approach, without encoding an anatomic model, to diagnose the cause of headaches on the basis of answers to two sets of questions, one set answered by the patient and the other by the physician. A positive response adds a predetermined number of points to one or more of 255 different types of headaches. If a requisite threshold level of points is reached for one type of headache, it is considered to be the diagnosis.

In general, there are two problems with approaches based on branching logic. First, their diagnostic capabilities are limited, because the diagnostic tasks to be fulfilled are poorly defined ("ill-structured"[6]), making it impossible for relatively simple-minded branching-logic algorithms to capture the expertise needed for the tasks. Second, numerical data structure and manipulation do not provide appropriate levels of abstraction for complex physical and logical modeling.

Approaches Based on Statistics

Statistics-based approaches try to arrive at a diagnosis by comparing data on an individual patient with data on a pertinent population (ie, obtained from large numbers of case records) that have been summarized numerically in a format suitable for calculating probabilities. These approaches are securely based on mathematical concepts and work well in small, established domains. Statistics-based approaches usually do not incorporate functional anatomic models or attempt to simulate the rational diagnostic strategies used by expert clinicians. Instead, on the basis of all Sxs, Pxs, and Lxs, they generally evaluate all possibilities by statistical methods and produce a rank order of the most likely diagnoses—a feature superficially resembling branching-logic evaluations of solution space.

Drawbacks of statistics-based approaches are that the computations involved may be extensive, because the searches are not focused, and that the amount of objective data required may be impossible to obtain. Because no cause-and-effect relationships are examined on the way to diagnostic conclusions, explanations take the form of "it is the case because it was the case." Conventional programming languages and statistics libraries can be used for statistics-based approaches to diagnosis.

Some investigators use statistical measurements of similarity. For example, a method designed to predict the site and type of intracerebral mass lesions[18] compares a new clinical database with 120 sets of data (Sxs, Pxs, radiologic findings, and histopathologically classified parenchymal lesions) from previous cases in order to calculate three levels of similarity indices. In essence, this is a form of statistical-pattern recognition in which the program tries to match a new patient's profile with templates (ie, patterns) of past cases, the best-matching template (according to certain criteria) providing the most likely diagnosis.

In one of its many medical applications, the famous Bayes's theorem[19] has been used to calculate the probability that a patient has a particular disease if he has a particular set of symptoms. According to the theorem, the probability of a disease, defined as its a posteriori probability, is proportional to the a priori probability of the disease (ie, the prevalence of the disease in the general population) multiplied by the probability of the symptoms (given the disease), defined as their likelihood (ie, the frequency with which the symptoms are associated with the disease). For likelihood of the symptoms, it is possible to substitute the product of the probability of each symptom, assuming that the symptoms are statistically independent of each other. The formulations used by many systems are more or less variants of the Bayesian inference scheme.[20-23] Symptom independence, nonconcurrence of diseases, and the need for vast amounts of statistical data about a general population are obstacles to the use of Bayes's theorem in computerized diagnosis. Interesting attempts to resolve the criticisms have been proposed.[24] In general, successful use of Bayes's theorem requires restriction of the characteristics of the patient population under study, thereby limiting the applicability of the results.

Broadly applicable epidemiologic data about a disease are difficult to obtain and are often unreliable. For these reasons, some investigators use the "maximum likelihood" method to quantify the relationship between a disease and a set of clinical data using frequency rates derived from a relatively small number of patients, thereby avoiding the pitfalls of population statistics. With these figures, the investigators identify that disease that maximizes the likelihood of the observed clinical database.[19] For example, in a program using brain scans for the differential diagnosis of brain lesions,[25] abnormal findings are analyzed in terms of 86 independent parameters; 240 proved cases are used to provide statistical data for the calculation of a diagnosis.

Approach Based on Techniques of Artificial Intelligence

Many early programs depended exclusively on branching logic or statistics, thus exemplifying the two extremes of the diagnostic decision-making spectrum[2]—categorical reasoning (eg, elimination of solution space) and probabilistic reasoning (eg, evaluation of solution space and statistical methods). There are strengths and weaknesses at both extremes. Categorical reasoning, which is particularly suitable when the logic of reasoning is well understood,[26-28] is efficient in problem solving. However, it is too rigid to handle complicated medical problems on which clinical data are incomplete or unreliable and for which clinching facts cannot be obtained. Probabilistic reasoning is useful when decisions must be made by carefully weighing all the evidence available, but it is practical only under certain assumptions and with a restricted solution space.

Every day, a physician faces a gamut of medical problems ranging from the extremely simple to the extremely complex. Problems at each extreme often require repetitive, sequential interplay between data acquisition and generation or confirmation of hypotheses to reach a diagnosis. The ideal diagnostic program must therefore be flexible and broadly applicable, and its diagnostic strategies must include both types of reasoning, regardless of the specialty involved. Conclusions such as these spawned a new generation of diagnostic programs, DESs.

The newer diagnostic systems employ some techniques of artificial intelligence, the most common being symbolic representation and pattern-directed inference, combined categorical and probabilistic reasoning, heuristic searching of the solution space, and generation and testing of hypotheses. These reflect the major investigative emphasis—computer simulation of intelligent human behavior—of artificial intelligence. Implementation of such systems relies heavily on LISP, the representative of symbolic processing programming languages.

Analogical Modeling

Neurologist, a DES dedicated to CNS diseases,[29] consists of four modules (*Input*, *Loc*, *Hgen*, and *Htest*) corresponding to the four sequential diagnostic steps: gathering the initial history and results of the physical examination, localizing the neurologic lesion, generating diagnostic hypotheses, and evaluating hypotheses. *Neurologist* maintains an analogical model of the parenchymal CNS that consists of 20 cross sections, five devoted to the spinal cord and the remainder to the brain stem and cerebral hemispheres. The cross sections are rough diagrammatic representations constructed from polylines (sequentially connected straight lines outlining boundaries) and convex polygons (each a polyline-enclosed figure in which a straight line between any two points falls entirely inside the figure). There are cross sections for about 100 nervous tracts (equivalent to axial neurosystems and transverse segment-limited structures in our terminology).

The system elicits patient information by means of predetermined "menus" displayed on a cathode ray tube and uses the answers to determine the status of nervous tracts according to simple logical criteria. Before it performs anatomic localization, the program matches and scores about 20 template descriptions against patterns of malfunction. Anatomic localization then consists of (a) attempts to find one or more cross sections permitting all the malfunctioning tracts to be included in a convex envelope[30] intersecting no more than one of the intact tracts and (b) the scoring of envelopes individually and jointly for adjacent cross section.[31,32] The most promising template or anatomic lesion, explaining over 80% of the findings, is used to localize the case; otherwise, the case is classified as nonlocalizable. Generation of differential diagnostic hypotheses from a table of 20 classes and a total of 120 "diseases" (in fact, a mixture of pathophysiological, pathogenetic, and disease-specific categories) is based primarily on the type of onset the illness exhibits and the localization. Each disease in the differential diagnosis is scored, and additional information is requested (using a hierarchical disease-attribute knowledge base).

Neurologist tries to mimic the diagnostic strategy of human experts (ie, rapid focus on a relatively small and promising part of the solution space for careful investigation), but it lacks the clinician's ability to shift attention when new data indicate possible causes outside the current considerations in differential diagnosis. Many other medical diagnostic systems, such as INTERNIST–1 and PIP, have paid great attention to this problem. Subtle balance of these seemingly contradictory actions is at the crux of a robust and efficient system.

The serial geometric cross-sectional representation of the CNS is a universal scheme used to model a three-dimensional complex. The larger the number of cross sections and the smaller the distance between adjacent cross sections, the more precise the model is. On the other hand, given the tasks the model is supposed to support, appropriate approximations can dramatically reduce the amount of data and the complexity of computation. Modeling the CNS by convex polygons allows, with relatively low computational complexity, geometric algorithms to be used as schemes for spatial reasoning. This modeling is an oversimplification, however, because the nervous tracts and anatomic lesions in each cross section usually need at least an arbitrary polygon (convex and concave) representation. Moreover, many anatomic concepts, including focality, seem to be captured better by topological propositions (eg, if all the malfunctioning tracts are adjacent to each other in a cross section then there may be a focal lesion in that cross section).[33]

Banks and Weimer[34] use the voxel approach to provide an anatomic knowledge base (SCAN) for neuroanatomic reasoning. They represent the human body as being embedded within a large cube. This cube is divided into 27 ($3 \times 3 \times 3$) small cubes, each of which is similarly subdivided into 27 smaller units. The process continues until the smallest cubes (currently 3 mm on each side) are reached. The resulting model is a hierarchy of nested cubes. The neuroanatomic components are represented by another hierarchy of "objects," each of which is mutually associated with its physical correspondent(s) in the cube hierarchy. While it is uniform and mathematically elegant, this approach is cognitively unnatural. The integrity of objects is not well preserved: an object represented by a single cube may also be represented by eight cubes of the same size, each of which involves part of the object merely because it is not aligned with the grid, thereby increasing the complexity of reasoning.

Propositional Modeling

First and associates[35] developed LOCALIZE, a system for identifying the sites of lesions in the PNS. It models the PNS in a propositional network in which muscles and spinal nerve roots are represented as terminal nodes, nerve segments as internal nodes, and anatomic connectivity relations as edges. Given a set of weak muscles on a scale from 0 (total paralysis) to 5 (normal strength), LOCALIZE traces the fibers that supply each affected muscle proximally toward the spinal cord and highlights the pathways. Any lesion or group of lesions including at least one highlighted segment from each traced pathway can hypothetically account for all deficits. Starting with the most distal set of lesions, the program generates alternative solutions by replacing lesion set elements with more proximal lesion sites from the highlighted pathways. The program tries to reduce the number of lesion sites hypothesized at each convergence point for as long as the consistency checks can be satisfied (ie, all muscles that the program expects to be weak because of a lesion at the convergence point are found to be affected, and normal muscles are not involved).

Goldberg and Kastner[36] have published a program designed to localize single, circumscribed (focal) lesions in the visual pathways. The program contains propositional, functional, and physical representations of the pathway and the adjacency relations between arbitrarily designated zones in each of 20 anatomically important cross sections through which nerve bundles pass from the retina to the cerebral cortex. Given a focal lesion, the program predicts the resulting defect in the visual field. Or, given information on a defect in the visual field, the system tries to identify a cross section in which a spatially continuous defect will encompass all involved bundles but none of the unaffected bundles. The reasoning depends on five heuristic, topological constraints intended to enforce the localizing principle of adjacency of damaged fibers. The program does not provide for the possibility of multiple lesions in the same or in different areas and neglects problems arising intrinsically in the central portions of a section—both common occurrences in neurologic disease.

These and other examples[37,38] demonstrate the power and usefulness of symbolic representation. The knowledge base captures effectively the most important structural information on underlying neuroanatomy, especially when the anatomic structure under consideration is no more than moderately complex—as is the case of the visual pathways. The knowledge base also provides for an association between structure and function on the one hand and spatially oriented reasoning on the other.

Rule-Based Architecture with No Explicit Anatomic Modeling

Rule-based DESs, the most elegant and successful developed thus far, have many medical and nonmedical applications. These DESs capture the human expert's tendency to employ judgmental knowledge to make decisions and, at the same time, satisfy such basic principles of software engineering as generality and modularity. MYCIN[2,39,40] is typical of a rule-based DES.

The basic representation of judgmental knowledge is simple and uniform: *if* premise *then* consequent *with* confidence. A rule in this format can be used for "forward chaining" and for "backward chaining" inference. In forward chaining inference, the premise tests the value of one set of parameters and the consequent contains conclusions about the value of another set of parameters. In backward chaining inference, the consequent matches the hypothesized value of one set of parameters while the premise verifies the value of another set of parameters. The application of one rule may trigger the application of other rules to form an inference chain. The confidence associated with each rule is used for "reasoning under uncertainty" (integrating categorical and probabilistic reasoning). Rules are stored in a "rule base," an independent component of the system.

The "inference engine," the control mechanism of this type of system, supervises the whole reasoning process. It initiates the inference activity, chooses the right rules, interprets the rules, and propagates the confidence information in a systematic manner (eg, using a modified Bayesian model,[41] the certainty factor theory exemplified by MYCIN, the theory of fuzzy sets,[39,42] or the Dempster-Shafer theory of evidence).[39,43]

From the point of view of reasoning, a rule-based system controls its attention by intermediate conclusions resulting from rules activated by current input—a natural match with the sequential nature of diagnosis. MYCIN uses a goal-directed, depth-first, backward-chaining control strategy in its inference engine to keep its focus during diagnosis. VM[44] employs a data-driven, forward-chaining strategy for cyclic reasoning in which previous observations and status-dependent expectations help direct progression from the clinical data to diagnostic and therapeutic decisions.

Creation of the rule base is the most important task in the construction of this type of DES, because the modular structure of rule-based architecture separates general reasoning mechanisms from domain-specific judgmental knowledge. Modification and expansion of the rule base are relatively easy. The explanation of a diagnosis can be obtained by keeping track of the rules used in the reasoning process. Dedicated tools supporting the development of such systems[45] include EMYCIN, EXPERT, OPS,[46] and PROLOG.[47]

Shortcomings of rule-based expert systems include the following: the need, when the rule base is very large, to introduce heuristic methods to index the most promising rules in any given context; an explanation capability which, in practical terms, is limited to dumping rules; and a lack of methodologies to integrate forward and backward inferences.

A DES for localizing CNS lesions in unconscious patients[48] has a rule-based architecture but lacks an explicit neuroanatomic model. Its rules map clinical findings to lesions at intracranial sites, and calculated values of certainty are used to differentiate the anatomic localization of clinically similar syndromes. The major drawback of this approach is that the system recognizes only encoded clinical patterns. Because it does not "understand" how the neurologic system is physically and functionally organized, this DES cannot reason anatomically when faced with unrecognized (unmatched) patterns.[48] This DES must therefore have access to an exhaustive list of rules covering anatomic interpretations for every conceivable combination of neurologic Sxs and Pxs—an impractical and unnatural requirement.

Comments

We conclude that little investigative attention has been devoted to the design of a knowledge representation effectively supporting a physical as well as a functional model of the entire neurologic system in the context of diagnostic reasoning. Perricone[7] and Reggia[49] discuss the importance of incorporating modeling schemes and reasoning mechanisms into DESs. To date, most of the proposed methods lack generality and do not support flexible reasoning mechanisms or a suitable interactive environment for the user.

We believe representational methods should accommodate both analogical and propositional information about neuroanatomy, support an association between structure and function, facilitate different levels of modeling, provide for diagnostic reasoning, and, finally, allow a variety of interfaces with the user.

NEUREX: TOWARD AN INTEGRATION OF MODELING AND REASONING

We turn now to our neurological DES, dubbed NEUREX (neurologic expert), which is evolving and experimental. After describing a general modeling scheme using a semantic network, we present details on our conceptual organization of the knowledge of neuroanatomic structure and function used to localize lesions and, finally, details on the representation of neuroanatomy.

Semantic Network Representation of Structure and Function

Cognitive knowledge is generally organized in the form of concepts and their relations to each other. A physical entity, each of its physical-spatial properties, and its function are all independent concepts that relate to each other when, in combination, they describe the entity. One can decompose a complex system into sets of entities, each corresponding to a different logical structure. The resulting sets of entities may or may not interweave with each other. The relations between entities (eg, spatial or functional connections) also are specific concepts.

A semantic network is a representation in which each concept (including each relationship between concepts) is represented by a specific "node" linked to other nodes by predefined "arcs."[50-53] Figure 4 shows the general organization of a semantic network representing spatial structure and function.

Others have used semantic network representations to implement expert systems. In PROSPECTOR, a geological analysis system,[54] knowledge is represented by partitioned semantic networks consisting of production rules and subset and element taxonomic information.

DESs in which reasoning is based on spatial structure and function gain several implementational advantages from a semantic network representation:

1. Analogical, propositional, and functional knowledge are integrated in a single network reflecting different levels of abstraction. Each physical entity is "surrounded" by its geometric and topological descriptions (if any), other spatially and/or functionally related entities, and its function. Thus a locally limited search of the network can provide all information relevant to the entity.
2. Rule-based inference is supported. A typical rule consists of two parts—antecedents and con-

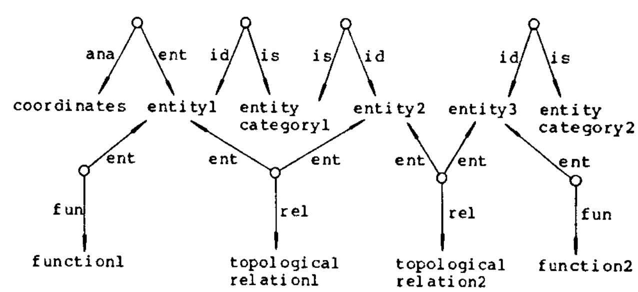

Figure 4. General organization of a semantic network representing spatial structure and function. Each conceptually significant entity is represented by a unique node, such as "entity1" or "entity2." The node with an "ent" arc pointing to entity1 and a "fun" arc pointing to "function1" asserts that entity1 has a function specified by function1. Analogical information about entity1 is asserted by the node with an "ana" arc pointing to "coordinates" and an "ent" arc pointing to entity1. Relations between entities, such as topological connections, are represented by nodes and arcs (eg, the nodes with "rel" and "ent" arcs). (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government*. Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

sequents—both of which may contain variables. To check whether the antecedent is satisfied, the program searches for the required nodes in the network. Consequents cause new nodes to be built. Complex control strategies can be superimposed on the basic network processing system.[54–57]

3. Easy expansion and modification provide flexibility. Knowledge is added or removed by the fundamental operations of adding concepts to or removing them from the network. Analogical data (coordinates) can be changed independently without changing any propositional information, provided that the relevant spatial relations continue to hold true.

4. Procedural knowledge is encoded in function nodes. For example, spatial reasoning often involves the application of algorithms from computational geometry using analogical data in the network.

5. Interactive graphics can be used for data entry and interpretive output. Analogical data can be used to generate relevant drawings, which, by means of locator or pick devices, can be used to enter data. Explanation capability also is greatly enhanced by the generation of appropriate drawings and pictures from analogical data.

6. New knowledge can be generated (and stored) from existing knowledge. For example, well-defined topological relations between physical entities can be computed systematically from existing analogical information and added to the network.

7. Computer vision techniques[8,9] can be supported. Both geometric and relational structures (derived from geometric descriptions and from topological relations embedded in the network), together with other relevant information, can be used to guide image segmentation and labeling and the interpretation of pictures.

8. A natural language interface between the diagnostic system and its users can be supported. The importance of semantic network representations in computational linguistics is well established.[58,59]

In summary, a semantic network is suitable for use in a model-based DES. Knowledge embedded in the network provides an understanding of how the subject is structured, how it works, and how it may fail.

Representation of Basic Knowledge in NEUREX

Knowledge concerning the spatial structure and function of neuroanatomy in NEUREX is organized as a semantic network according to the rationale presented in the preceding section. Representation is implemented in the FRANZ LISP system (a LISP dialect developed at the University of California at Berkeley)[60] and the Semantic Network Processing System (SNePS)[51] on a VAX-11/750 computer under the 4.2 BSD UNIX operating system. This LISP system and SNePS are completely compatible and mutually callable. SNePS allows effective searches of the network knowledge base and the incorporation of both procedural attachments (function nodes) and inference rules.

The model described in the following pages is being used to study problems in the design and implementation of our approach. At present, we have encoded information about the entire spinal cord and the arms, along with relevant graphic reconstructions.

The anatomic knowledge base contains information about neuroanatomy and general anatomy. The CNS is divided into its major transverse segments: telencephalon, diencephalon, brain stem (mesencephalon, pons, and medulla), and spinal cord, through which neurologic data are transmitted along multiple axial neurosystems and in which these data are analyzed, integrated, relayed, and supplemented or removed. Each major transverse segment is subdivided into smaller, naturally occurring or arbitrarily defined segments and regions within segments, to facilitate precise localization. If appropriate, a transverse segment connects on its right and left sides with the PNS via PNS nerve roots or their equivalents (ie, cranial nerves attached to the brain stem). A PNS root is not identified by the CNS segment to which it is attached. Each root network (CNS–PNS) innervates (transmits multiple types of encoded sensory information from and/or motor directives to) specific regions or structures of the body. Except for the side of origin and peripheral distribution, homologous right and left CNS–PNSs are anatomically identical and innervate corresponding areas of the body. Each CNS–PNS passes through a system of conduits or peripheral nerves to reach its termini. As they extend from the CNS, peripheral nerves assume a branch-tree structure that segments them "transversely," first into spinal or cranial nerves and then into as many additional subunits as are needed to reach the final nerve segments. Thus the human neurologic system consists of many axial neurosystems, each of which has a CNS component. Some are limited to the CNS; others (eg, the somatic motor and sensory neurosystems) also extend into the PNS. Therefore, the latter neurosystems have three general patterns of innervation: one corresponding to CNS pathways; the second, to CNS–PNSs; and the third, to peripheral nerves.

The major functional axial neurosystems transmitting information up or down the neurologic system are made up of smaller axial units. Each of these mini-neurosystems is identified uniquely by (1) the specific areas it occupies as it passes through each transverse segment of the CNS or the CNS and PNS; (2) the transverse segment of the CNS in which it originates, terminates, and/or connects with the

PNS; (3) the direction in which it carries data; and (4) a clinically verifiable function. A functional mini-neurosystem represents a large number of individual nerve fibers relaying the same information in parallel and in series. The point at which one set of its fibers connects with the next is indicated in the model, even though the transferred function does not change.

We are currently working on CNS, CNS–PNS, and peripheral nerve representations of the axial somatic sensory/somatesthetic and somatic motor neurosystems, particularly as they relate to the spinal cord. Both neurosystems have central and peripheral components. The mini-neurosystems responsible for somatesthetic information destined to reach consciousness transmit from the lowest segments of the spinal cord to higher segments of the brain ("upward"). The major volitional motor mini-neurosystems transmit directives caudally ("downward"). Regardless of the direction in which they conduct (and with exceptions related to intersegmental reflexes), the number of mini-neurosystems in the major axial pathways decreases in a rostral-to-caudal direction in relation to other segment-limited structures or to peripheral nerves. (Sensory mini-neurosystems enter from the peripheral nerves and ascend predominantly, while motor mini-neurosystems originating rostrally descend to terminate in succeeding transverse segments.) Therefore, at any given point along the central neuraxis, a CNS pattern of motor or sensory innervation will depend on the number of mini-neurosystems in the axial pathway; the pattern involves the entire field of innervation to the half of the body supplied by the remaining mini-neurosystems. CNS–PNS ("cranial nerve," and "spinal root" or "segmental") patterns correspond to the peripheral special sensory, somatosensory, and motor innervation of the cranial nerve or spinal root. The following discussion is limited to CNS–PNSs related to the spinal cord.

With rare exception, each spinal CNS–PNS carries several somatic sensory and two types of somatic motor mini-neurosystems. A peripheral nerve sensory or motor pattern may or may not be identical to that of a single CNS–PNS. Those related to the limbs are complex. Several adjacent CNS–PNSs combine in a complicated manner at specific junctional points (plexi) and then partially dissociate at more distal branch points. Therefore, a proximal peripheral nerve may involve sensory and/or motor mini-neurosystems (and innervation fields) from two or more CNS–PNSs. The composition of its more distal segments will be the same or less complex, depending on partial dissociation of the mini-neurosystems at any succeeding branch point. Conversely, a CNS–PNS may traverse many or relatively few peripheral nerves. Despite the complexity of the CNS, CNS–PNS, and peripheral nerve patterns, the anatomic pathways involved are unique from origin to final destination and very consistent from one individual to the next.

In the knowledge base, the body is divided into its major regions (eg, head, arms, torso) and subregions (eg, upper, middle, and lower third of the right brachium; shoulder, elbow, wrist, interphalangeal joints). The subregions serve as anatomic reference points, helping to locate the components of the neurologic system, supporting maps of the cutaneous distribution of sensation, and organizing skeletal muscle function.

Structural features of the components of the nervous system are attached to their corresponding concepts in different ways, according to the significance of the features: coordinates of the polyline and polygon representation are used for cross sections, regions, and projected views of body parts; connectivity relationships are used for nerves and branches of nerves. All the data are accommodated in a single information base.

Representing the CNS, CNS–PNSs, Peripheral Nerves, and Their Functional Distribution

The CNS can be described geometrically by cross sections through transverse segments. Every cross section and every region in the section is an anatomic concept that, among other properties, has a geometric description (ie, unique coordinates in a common coordinate system). More abstract relations, such as the adjacency of two regions, can be asserted between the corresponding concepts. Meanwhile, each axial or transverse neuroanatomic pathway is a concept defined further by a set of otherwise independent concepts (for example, assignment to an axial neurosystem, spatial characteristics, anatomic locations, specific function, evidence of malfunction). This is shown in Figure 5, a computer-generated, schematic reconstruction of the clinically significant axial neurosystems and segment-limited structures in the fifth cervical segment of the spinal cord.

The CNS–PNSs and peripheral nerves are represented topologically as a network in which each transverse segment is an anatomic concept. Connectivity of segments is asserted between corresponding concepts. When a particular axial mini-neurosystem travels through the network, its pathway is specified by assertions relating the mini-neurosystem to the segments through which it passes. Each mini-neurosystem is identified by its parent neurosystem, its specific function(s), and its peripheral innervation field.

To represent patterns of CNS–PNS and peripheral nerve innervation, we consider each region or structure of the body to be an anatomic concept, each with a geometric description outlining a corresponding region in a display (Figure 6). Information such as the transverse nerve segments or CNS–PNS supplying an area is attached by assertions.

Therefore, we intend to represent the structural and functional knowledge necessary for neurologic diagnosis in a semantic network in which anatomic concepts are fundamental entities.

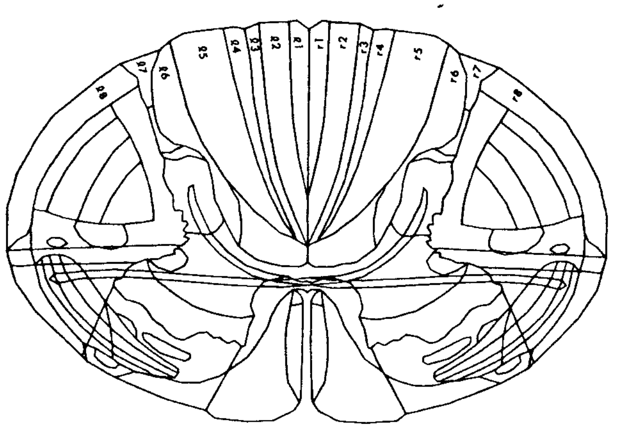

Figure 5. Schematic of clinically significant axial neurosystems and segment-limited structures in the fifth cervical segment of the spinal cord. Each labeled area represents an anatomic region through which fibers of one or more axial neurosystems pass as they intersect the transverse section. The right fasciculus gracilis, carrying several types of sensory neurosystem/somatesthetic fibers, intersects the fifth cervical segment in regions r1, r2, and r3. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

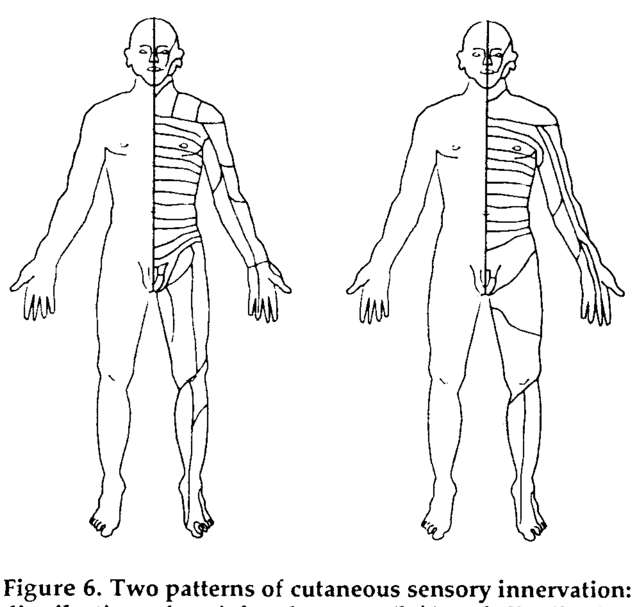

Figure 6. Two patterns of cutaneous sensory innervation: distribution of peripheral nerves *(left)* and distribution of transverse segments of the spinal cord and their attached peripheral root systems (CNS–PNS pattern) *(right)*.

Detailed Representation of Neuroanatomy

The first set of anatomic concepts (cross sections and regions in each cross section in the CNS, each CNS–PNS, each peripheral nerve, and regions in the distribution of peripheral nerves and CNS–PNSs) is represented by atomic nodes. Each atomic node may have a geometric description—a sequence of coordinates that can be linked to form polylines or polygons in a given coordinate system. Abstract spatial relations are asserted between nodes.

For example, the fifth cervical segment of the spinal cord in Figure 5 is represented by atomic node C5 in Figure 7. Node C5 has a geometric description carrying the coordinates of its boundary. Each region in the cross section is represented by an atomic node (eg, r1 in Figure 5 is represented by node r1C5 in Figure 7). Node r1C5 has a geometric description carrying the coordinates of its boundary. That r1 and r2 are adjacent to each other in C5 is asserted by the node with "r" arcs pointing to r1C5 and r2C5 and a "rel" arc pointing to "adjacent" in Figure 7. Other relations, such as overlap, also can be asserted.

Another set of concepts covers the major axial neurosystems ("tracts" or "pathways") in the CNS. Each tract is represented by an atomic node. For example, node 7R in Figure 7 represents the right fasciculus gracilis. The physical location of the tract in C5 is specified by an assertion indicating the corresponding regions by "path" arcs.

Anatomically significant components of the CNS–PNS and transverse nerve segments of the PNS are represented by unique atomic nodes, and the connectivity relations are specified by nodes with proxi-

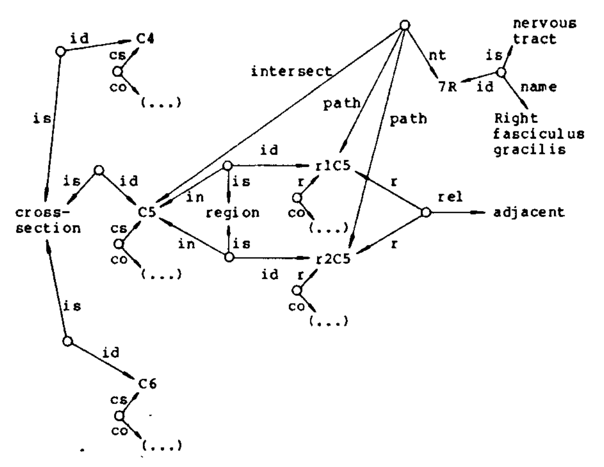

Figure 7. A semantic network partially representing the fifth cervical transverse segment of the spinal cord (C5). Spatial relations and coordinates encoded in atomic nodes and arcs translate graphically into polylines and polygons. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

mal and distal arcs. Figure 8 illustrates representation of the simple CNS–PNS system shown in Figure 3. The pathway of a particular mini-neurosystem can be traced by a sequence of nerve segments and branches linked by *p* and *d* arcs from its origin to its destination.

In the neuroanatomic model built thus far, additional information concerning anatomy and function will be represented as concepts and inserted into the network according to the same principle. The somatic motor neurosystem, for instance, is responsible for movement of joints. The movement of a joint is controlled by several muscles, and muscles may be supplied by motor mini-neurosystems related to one or, as in the case of the limbs, more than one CNS–PNS. If more than one CNS–PNS innervates a muscle, the motor mini-neurosystem carried by each CNS–PNS contributes a certain percentage to the total innervation of the muscle and, therefore, to regulating the force of the muscle's contraction. Furthermore, because a particular movement at a joint may require the synchronous contraction of several muscles, a muscle may be responsible for all or only a fraction of the force of a movement. Muscles, joints, different types of movement, the contributions of each muscle to a movement, and the contributions of each peripheral nerve and each CNS–PNS to the innervation of a muscle are all represented by atomic nodes between which functional relations are asserted.

For example, the shoulder joint (movement of the humerus in relation to the scapula) has eight different types of movement: flexion (forward), extension (backward), abduction 0 to 15 degrees, abduction 15 to 90 degrees, adduction, external (lateral) rotation, internal (medial) rotation, and rotator cuff. Flexion is controlled by the following muscles: deltoid (40%), pectoralis major–clavicular head (45%), coracobrachialis (10%), and biceps brachii (5%); extension by deltoid (30%), teres major (20%), and latissimus dorsi (50%); abduction 0 to 15 degrees by supraspinatus (90%) and biceps brachii (10%); and abduction 15 to 90 degrees by deltoid (100%). The percentage in each set of parentheses is an arbitrary approximation of the contribution of the particular muscle to the total strength of the movement. Figure 9 illustrates the representation of this information.

Motor mini-neurosystems from each CNS–PNS innervating the muscle are treated in a similar manner. The complete pathway of a mini-neurosystem passing through the CNS and PNS is represented by other nodes, specifying its peripheral nerve pathway and the anatomic regions of each transverse CNS segment through which it travels. The anatomy of the PNS is not depicted in graphics comparable to those used for the CNS, because the level of abstraction we are using for the PNS favors anatomic localization reasoning but does not provide geometric details for graphics display.

Additional Details and Examples

The following details and examples demonstrate each of the advantages of semantic network representation stated above:

1. A request for the location of a particular mini-neurosystem in a cross section of the CNS produces a locally limited search of the network, gen-

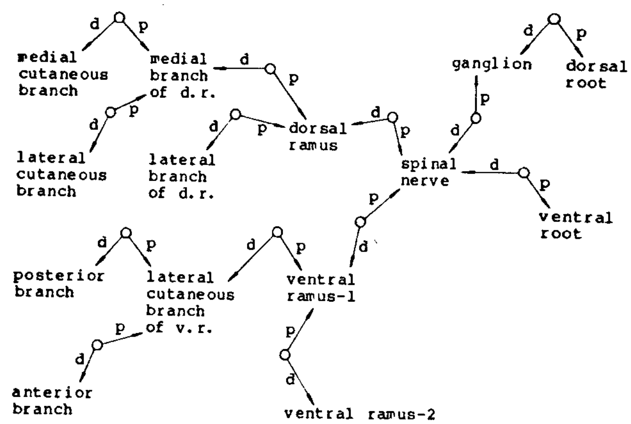

Figure 8. A semantic network partially representing the single peripheral spinal root system shown in Figure 3. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

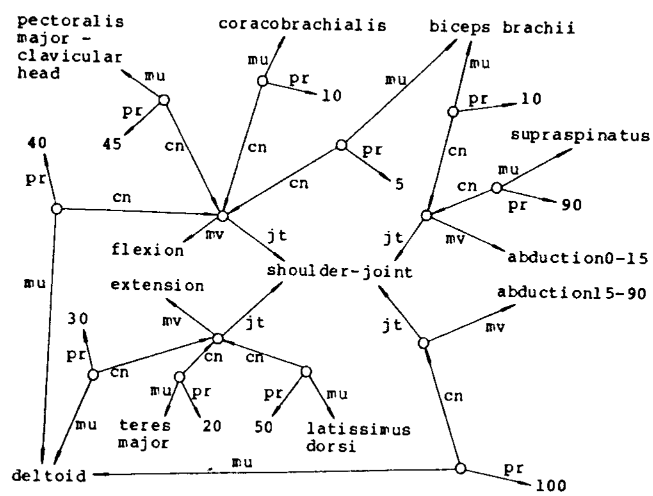

Figure 9. A semantic network partially representing functional muscular anatomy of the shoulder joint. The nodes and arcs can describe the eight cardinal movements of the joint, the muscles involved, and the contribution of each muscle to each movement. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

erates a picture such as Figure 5 on the screen, and in the picture highlights the region and lists all other mini-neurosystems passing through the same area (along with their function, the major axial neurosystem to which each belongs, their origins and terminations, the direction in which they conduct information, and so on).

Figure 9 shows the use of the uniform network-searching function in SNePS. To find out which muscles are involved in flexion of the shoulder joint, we issue the request (find (mu- cn) (find jt shoulder-joint mv flexion)), which returns a list containing the names of the four muscles involved. To find out how much the deltoid contributes to the strength of shoulder flexion, we issue the request (find pr- (find mu deltoid cn (find jt shoulder-joint mv flexion))) and the program returns: (40). To find out how paralysis of the deltoid affects shoulder movement, we issue the request (find mv- (find jt shoulder-joint (cn- mu) deltoid)) and the program returns: (flexion extension abduction 15-90).

2. A rule might be stated as follows: if a particular mini-neurosystem carried by a certain peripheral-nerve segment is malfunctioning, then other mini-neurosystems in the same segment must be examined for malfunction. This rule is displayed graphically in Figure 10. Rules can include probabilistic data, such as certainty factors, and we can perform the necessary computations by procedural attachment.

3. The information base can be revised without difficulty or disruption of other relationships in the network. If, for instance, research discloses that a mini-neurosystem passes through a region in a cross section of the CNS different from that currently specified in the network, a new node

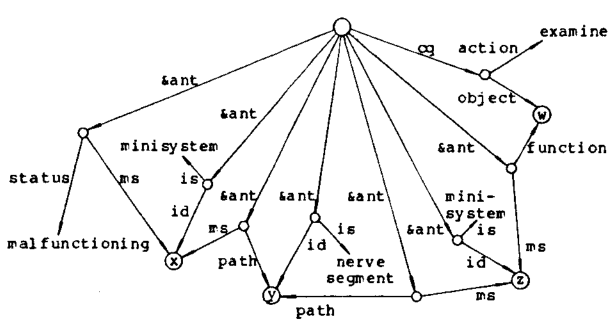

Figure 10. A rule for neuroanatomic localization represented in a semantic network. The rule states that if a mini-neurosystem carried by a specific peripheral-nerve segment is malfunctioning, then other mini-neurosystems in the same segment must be examined for malfunction. (Published originally in Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. A modeling scheme for diagnosis. In: Karna KN, ed. *Expert systems in government.* Washington, DC: IEEE Computer Society, 1985:538-547. © 1985 IEEE.)

specifying the relation between the new region and the mini-neurosystem can be added, and the old node relating the original region and the mini-neurosystem can be removed.

4. SNePS supports diagnostic reasoning about anatomy (neuroanatomic localization). Sxs, Pxs, and Lxs are mapped to specific axial neurosystems in a given transverse segment; a closed curve encompassing all the malfunctioning mini-neurosystems and no normally functioning structure in the area is defined. Or, as another example, because several adjacent or overlapping abnormal regions in the same cross section tend to be affected by a single lesion, the task of determining the minimum number of lesions, given a set of abnormal regions, is, in fact, the task of determining the number of spanning trees in an undirected graph in which a vertex represents a region and an edge represents an adjacency or overlap relation.

5. SNePS supports graphics. The geometric information about CNS–PNS and peripheral-nerve patterns is used to display pictures such as Figures 1 and 6 on the screen. By means of a stylus or similar device, the extent of a sensory Sx or Px can be indicated on the picture. Localization begins by storing the picture in an image array, counting the number of pixels involved in different regions, and comparing these to the appropriate patterns of CNS–PNS and peripheral-nerve innervation to establish which of the latter most closely approximates the extent of the lesion. The results are stored as new nodes for further reasoning.

For example, the lesion might involve complete loss (100%) of light touch in the area about the chin and lower lip, as shown in Figure 1 (right). If one asks which CNS, CNS–PNS, or peripheral-nerve pattern best describes the lesion, the system returns the name of the peripheral nerve to the chin, rejecting all other possibilities as unlikely. On the other hand, if the left side of the body is involved from the level of the nipple downward, the system returns the name and a drawing of the most rostral (transverse) segment of the spinal cord containing all the somatesthetic–light touch mini-neurosystems involved and none of those uninvolved, along with an outline of the lesion. If the lesion extends through several transverse segments, the geometric information can be used to construct a three-dimensional picture on the screen. The directive can also be reversed, for example, by drawing a lesion on the display of a particular transverse segment of the CNS and asking the system to outline the expected sensory loss on the appropriate graphic display of the body. Although somatic sensory neurosystems were used to demonstrate some features of localization, the same principles can be demonstrated just as well with somatic motor neurosystems, the input or output being the weakness of muscles.

6. Drawings such as those in Figures 1, 5, and 6 can be created with a geometric graph editor.[61] The adjacency relationship is inserted into the knowledge base by a system-construction function that asserts, for every cross section, an adjacency relation between each pair of regions sharing a common boundary.
7. For computer-assisted tomography of the CNS, geometric information in the knowledge base can be increased to provide realistic, age-dependent, geometric structures corrected statistically for normal variation. The system could then be expanded to interface with scanners (roentgenographic and nuclear magnetic resonance) to relate tomographic output to functional neuroanatomy.
8. A natural language interface for literal transactions (eg, patient data entry and explanation of diagnosis) can be developed.

CONCLUSIONS

The construction of expert systems is a branch of computer science that relies on philosophical and psychological concepts to support or refute its various experimental approaches to the simulation of intelligent human behavior. The history of computerized neurological diagnosis outlined in this paper exemplifies the development of DESs in general. Time and effort devoted to this area of research will probably enhance our insight into human cognition and lead to the development of more general methodologies and implementational tools applicable to other domain-specific systems.

As we try to integrate modeling and reasoning, we confront many fundamental issues, ranging from the representation of knowledge to the mechanisms of reasoning. The central task of "knowledge engineers" is to find the point at which machines perform intelligently as well as or better than humans—a point at which the knowledge needed to perform the task is represented adequately and economically, the reasoning with computers simulates or conceivably improves upon the human approach, and the tools needed for implementation are practical and user-friendly.[62,63] As a general method in which to embed the fundamental expertise of model-based DESs, semantic networks are flexible enough to handle different kinds of information, to support diagnostic reasoning, to allow for natural interfaces with the user, and to permit learning. Semantic networks are also sufficiently simple and uniform to facilitate the development of processing tools.

REFERENCES

1. Winston PH. *Artificial intelligence.* 2nd ed. Reading, MA: Addison-Wesley, 1984.
2. Clancey WJ, Shortliffe EH, eds. *Readings in medical artificial intelligence: The first decade.* Reading, MA: Addison-Wesley, 1984.
3. Chandrasekaran B, Mittal S. Deep versus compiled knowledge approaches to diagnostic problem-solving. In: Coombs MJ, ed. *Developments in expert system.* New York: Academic Press, 1984:23-34.
4. Davis R, Shrobe H, Hamscher W, et al. Diagnosis based on description of structure and function. Presented at the National Conference on Artificial Intelligence, Pittsburgh, 1982.
5. Davis R. Diagnostic reasoning based on structure and behavior. *Artificial Intell* 1984; 24:347-410.
6. Pople H. Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnosis. In: Szolovits P, ed. *Artificial intelligence in medicine.* American Association for the Advancement of Science, Selected Symposium 51. Boulder, CO: Westview Press, 1982:119-190.
7. Perricone BT. Spatial representation and reasoning: A critical survey and proposal, in representation and processing of spatial knowledge. TR-1275, Department of Computer Science. College Park, MD: University of Maryland, May 1983.
8. Srihari SN. Representation of three-dimensional digital images. *Comput Surv* 1981; 13:399-424.
9. Ballard DH, Brown CM. *Computer vision.* Englewood Cliffs, NJ: Prentice-Hall, 1982.
10. Adams R, Victor M. *Principles of neurology.* New York: McGraw-Hill, 1985.
11. Chutkow JG. Neurology manual. Buffalo: Department of Neurology, School of Medicine and Affiliated Hospitals, State University of New York at Buffalo, 1984.
12. Freemon FR. Computer diagnosis of headache. *Headache* July 1968; 8:49-56.
13. Vastola EF, Thomasino J, Steger L, Lindsay W. Evaluation of a computer program for neurological screening. *Bio-Med Comput* 1973; 4:117-129.
14. Hofferberth B, Gottschaldt M. The writing of a case history in a neurological department with computer application. *Methods Inf Med* 1978; 17(3):162-166.
15. Ropper AH, Griswold K, McKenna D, Souder D. Computer-guided neurologic assessment in the neurologic intensive care unit. *Heart Lung* 1981; 10(1):54-60.
16. Meyer A, Weissman W. Localization of lesions in the human nervous system by computer analysis. *IEEE Trans Bio-Med Eng* 1973; 20:194-200.
17. Francis JH, Pennal BE, Wadsworth W. Development of a computer-assisted headache diagnosis and treatment system. *Headache* 1984; 24:35-38.
18. Stewart A, Cala L. Mathematical method to utilize a computer for diagnosis of site and type of intracerebral mass lesions. *Br J Radiol* 1975; 48:97-100.
19. Kendall MG, Stuart A. *The advanced theory of statistics 1.* New York: Hafner, 1963:198-204.
20. Croft DJ. Is computerized diagnosis possible? *Comput Biomed Res* 1972; 5:351-367.
21. Gorry GA, Barnett GO. Experience with a model of sequential diagnosis. *Comput Biomed Res* 1968; 1:490-507.
22. Gorry GA, Kassirer JP, Essig A, Schwartz WB. Decision analysis as the basis for computer-aided management of acute renal failure. *Am J Med* 1973; 55:473-484.

23. De Boulay GH, Harling D, Clarke G. Improvement in the computer-assisted diagnosis of cerebral tumours. *Br J Radiol* 1977; 50:849-854.
24. Charniak E. The Bayesian basis of common sense medical diagnosis. *Proc Am Assoc Artificial Intell* 1983; 70-73.
25. Mori H, Suzuki Y, Hisada K, et al. Differential diagnosis of brain lesions with a computed brain scan diagnosis by the likelihood method. *Invest Radiol* 1975; 10:251-257.
26. Feinstein AR. An analysis of diagnostic reasoning. I. The domains and disorders of clinical macrobiology. *Yale J Biol Med* 1973; 46:212-232.
27. Feinstein AR. An analysis of diagnostic reasoning. II. The strategy of intermediate decisions. *Yale J Biol Med* 1973; 46:264-283.
28. Feinstein AR. An analysis of diagnostic reasoning. III. The construction of clinical algorithms. *Yale J Biol Med* 1974; 47:5-32.
29. Catanzarite VA. *Neurologist*: A computer program for localization and diagnosis in clinical neurology. Ph.D. diss, Department of Biophysics, University of California, Berkeley, 1980.
30. Preparata FP, Hong SJ. Convex hulls of finite sets of points in two and three dimensions. *Commun ACM* 1977; 20(2):87-93.
31. Bentley JL, Ottmann TA. Algorithms for reporting and counting geometric intersections. *IEEE Trans Comput* 1979; C-28(9):643-647.
32. Shamos MI, Hoey D. Geometric intersection problems. In: *17th Annual Symposium on Foundations of Computer Science*. New York: IEEE, 1976:208-215.
33. Xiang Z, Srihari SN, Shapiro SC, Chutkow JG. Analogical and propositional representations of structure in neurological diagnosis. Presented at the First Conference on Artificial Intelligence Applications, Denver, 1984.
34. Banks G, Weimer B. Symbolic coordinate anatomy for neurology (SCAN). *J Med Syst* 1984; 8(3):157-162.
35. First M, Weimer B, McLinden S, Miller R. LOCALIZE: Computer-assisted localization of peripheral nervous system lesions. *Comput Biomed Res* 1982; 15:525-543.
36. Goldberg R, Kastner J. An explicit description of anatomical knowledge as an aid to diagnosis: A spatial network model. CBM-TM-78. New Brunswick, NJ: Rutgers University Department of Computer Science, 1978.
37. Chen JK. Representation of neurological paths by semantic networks. M.S. project report, Department of Computer Science, State University of New York at Buffalo, 1984.
38. Suchin J. A semantic network representation of the peripheral nervous system. M.S. project report, Department of Computer Science, State University of New York at Buffalo, 1985.
39. Buchanan BG, Shortliffe EH, eds. *Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley, 1984.
40. Cendrowska J, Bramer M. Inside an expert system: A rational reconstruction of the MYCIN consultation system. In: O'Shea T, Eisenstadt M, eds. *Artificial intelligence: Tools, techniques, and applications*. Hagerstown, MD: Harper & Row, 1984:453-497.
41. Duda RO, Hart PE, Nilsson HJ. Subjective Bayesian methods for rule-based inference systems. *Proceedings of the 1976 AFIPS National Computer Conference*. 1976:1075-1082.
42. Orci IP. Programming in fuzzy logic for expert systems design. *Proceedings of the Fifth International Workshop on Expert Systems and Their Applications*. Avignon: Agence de l'Informatique, 1985:1179-1192.
43. Oblow EM. A hybrid-uncertainty theory. *Proceedings of the Fifth International Workshop on Expert Systems and Their Applications*. Avignon: Agence de l'Informatique, 1985:1193-1201.
44. Fagan LM, Shortliffe EH, Buchanan BG. Computer-based medical decision making: From MYCIN to VM. In: Clancey WJ, Shortliffe EH, eds. *Readings in medical artificial intelligence: The first decade*. Reading, MA: Addison-Wesley, 1984:241-255.
45. Hayes-Roth F, Waterman DA, Lenat DB, eds. *Building expert systems*. Reading, MA: Addison-Wesley, 1983.
46. Forgy C, McDermott J. OPS: A domain-independent production system language. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Los Altos, CA: William Kaufman, Inc., 1977:933-939.
47. Clocksin W, Mellish C. *Programming in PROLOG*. New York: Springer-Verlag, 1981.
48. Reggia JA. A production rule system for neurological localization. Presented at the 2nd Annual Symposium on Computer Applications in Medical Care, Washington, DC, 1978.
49. Reggia JA. The role of spatial knowledge in expert systems, in representation and processing of spatial knowledge. TR-1275, Department of Computer Science. College Park, MD: University of Maryland, May 1983.
50. Hendrix GG. Encoding knowledge in partitioned networks. In: Findler NV, ed. *Associative Networks: Representation and use of knowledge by computers*. New York: Academic Press, 1979:51-92.
51. Shapiro SC. The SNePS semantic network processing system. In: Findler NV, ed. *Associative networks: Representation and use of knowledge by computers*. New York: Academic Press, 1979:121-175.
52. Shapiro SC. A net structure for semantic information storage, deduction and retrieval. *Proceedings of the Second International Joint Conference on Artificial Intelligence*. Los Altos, CA: William Kaufman, Inc. 1971:368-374.
53. Woods WA. What's in a link: Foundations for semantic networks. In: Bobrow DG, Collins AM, eds. *Representation and understanding*. New York: Academic Press, 1975:35-82.
54. Duda RO, Hart PE, Nilsson NJ, Sutherland GL. Semantic network representations in rule-based inference systems. In: Waterman DA, Hayes-Roth F, eds. *Pattern-directed inference systems*. New York: Academic Press, 1978:203-221.
55. Martins JP, Shapiro SC. Reasoning in multiple belief spaces. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*. Los Altos, CA: William Kaufman, Inc., 1983:370-373.
56. McKay DP, Shapiro SC. Using active connection graphs for reasoning with recursive rules. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*. Los Altos, CA: William Kaufman, Inc., 1981:368-374.
57. Shapiro SC. Path-based and node-based inference in

semantic networks. In: Waltz D, ed. *TINLAP–2: Theoretical issues in natural language processing–2*. New York: ACM, 1978:219-225.

58. Shapiro SC. Generalized augmented transition network grammars for generation from semantic networks. *Am J Computat Linguist* 1982; 8(1):12-25.
59. Sowa J. *Conceptual structures: Information processing in mind and machine.* Reading, MA: Addison-Wesley, 1984.
60. Foderaro JK, Sklower KL, Layer K. The FRANZ LISP Manual. In: *UNIX programmer's manual*, Suppl doc 4.2. Berkeley Software Distribution, Virtual VAX-11 version. Berkeley: University of California, Berkeley, Department of Electrical Engineering and Computer Science, 1984.
61. Xiang Z, Srihari SN. Graphics interfaces to a neurological diagnosis system. Presented at the Sixth Annual Conference and Exposition 3 of the National Computer Graphics Association, Dallas, 1985.
62. Miller RA. INTERNIST–1/CADUCEUS: Problems facing expert consultant programs. *Methods Inf Med* 1984; 23(1):9-14.
63. Kulikowski CA. Artificial intelligence methods for expert medical consultant systems. *Mt Sinai J Med* 1985; 52(2):87-93.

EDITORIAL COMMENT

Zhigang Xiang, who received his B.S. in computer science from Beijing Polytechnic University, is a doctoral candidate in computer science at the State University of New York at Buffalo. His research interests include expert systems, computer system architecture, and parallel processing.

Jerry G. Chutkow, M.D., professor and past chairman of the Department of Neurology at the School of Medicine, State University of New York at Buffalo, received his M.D. degree from the University of Chicago. Following residencies in internal medicine and neurology at the University of Chicago Hospitals and Clinics, he served on the faculty of the Mayo Medical School, Rochester, Minnesota, before joining the faculty at SUNY Buffalo in 1977. His research interests include the metabolism of cations in the central nervous system, neuromuscular disease, and computer-assisted diagnosis in neurology.

Dr. Shapiro is professor and chairman of the Department of Computer Science, and Dr. Srihari is an associate professor of computer science, at the State University of New York at Buffalo.

Michael Anbar
Editor-in-Chief