

□ USING CAUSAL REASONING IN GAIT ANALYSIS

DAVID E. HIRSCH

Department of Electrical Engineering and Computer
Science, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139

SHELDON R. SIMON

Department of Surgery, The Ohio State University,
Columbus, Ohio 43210

TOM BYLANDER and MICHAEL A. WEINTRAUB

Department of Computer and Information Science,
The Ohio State University, Columbus, Ohio 43210

PETER SZOLOVITS

Department of Electrical Engineering and Computer
Science, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139

This paper describes a series of experiments in which expert diagnostic systems were constructed to analyze human pathologic gait. The difference between successive systems is the recognition of the need for both causal reasoning about the process of gait and experiential, associational knowledge that can control causal reasoning. The performance of the first system (DR. GAIT-1), which relies exclusively on associational knowledge, is quite limited. The second system (DR. GAIT-2), because it is based on a qualitative causal model of gait, overcome many of the difficulties faced by the first system, but its ability to diagnose cases is limited by the complexity of causal reasoning. The third system (QUAWDS), which we are currently developing, is an experiment in integrating causal reasoning with associational knowledge so that robust conclusions can be produced efficiently.

INTRODUCTION

Diagnosis is the task of explaining a set of observations in terms of malfunctions and their causes. This paper describes a series of experiments in which expert diagnostic systems were constructed to analyze human pathologic gait. The experiments consist of three different systems, all utilizing the knowledge

This research has been funded by contract G008300074 from the National Institute for Handicap Research, grant R24 RR01320 from the Division of Research Resources of the National Institutes of Health, and grants 82048-02 and H133E80017 from the National Institute on Disability and Rehabilitation Research. Computer facilities were enhanced through a grant from Digital Equipment Corporation and gifts from Xerox Corporation.

David E. Hirsch is now at Price Waterhouse Technology Center, Menlo Park, California.

Correspondence should be sent to Tom Bylander, Computer and Information Science, Ohio State University, Columbus, OH 43210.

provided by one of the authors (Simon). The difference between successive systems is the recognition of the need for both causal reasoning about the process of gait and experiential, associational knowledge that can control the causal reasoning.

The first system (DR. GAIT-1), developed by three of the authors (Hirsch, Simon, and Szolovits) at MIT and Harvard, relies exclusively on associational knowledge. DR. GAIT-1's success in diagnosing cases is quite limited. The second system (DR. GAIT-2), developed by the same three people, is primarily based on a qualitative causal model of gait. DR. GAIT-2 overcame many of the difficulties faced by the first system, but its ability to diagnose cases is limited by complexity of causal reasoning. The third system (QUAWDS, for QUalitative Analysis of Walking Disorders), currently being developed by three of the authors (Bylander, Weintraub, and Simon) at The Ohio State University, is an experiment in integrating causal reasoning with associational knowledge so that robust conclusions can be produced efficiently.

First, we briefly describe the domain of gait analysis. Next, we discuss each system in turn, with special attention given to the role of causal reasoning within each system. Because of space limitations, our descriptions are necessarily brief and simplified. Also, our attention is mainly focused on the diagnostic functions of these systems although they also provide recommendations for therapy. For further information on DR. GAIT-1 and DR. GAIT-2, see Hirsch (1987). For further information on QUAWDS, see Bylander et al. (1988).

THE DOMAIN OF HUMAN PATHOLOGIC GAIT

Normal gait is efficient, adaptable, pain-free, and requires no ancillary devices. In a normal person, the neurological system controls the muscles through coordinated commands to rotate limbs at several joints, providing body propulsion and stability for walking (Inman et al., 1981; Perry, 1985). A gait cycle consists of the time between a heel strike and the next heel strike of the same foot. The most significant events of the gait cycle are right heel strike (RHS), left toe off (LTO), left heel strike (LHS), and right toe off (RTO), which delimit the major phases of gait: weight acceptance (WA), single limb stance (SLS), weight release (WR), and swing. Figure 1 illustrates these events and the phases for the right leg.

The domain of DR. GAIT-1 and DR. GAIT-2 is restricted to pathologic gait resulting from the neurological disease of cerebral palsy. This disease affects the brain and manifests itself by interfering with the control of voluntary and involuntary motions. The effect of cerebral palsy on the gait cycle is improper coordination of muscle activity. It is these effects, and not cerebral palsy itself, that are the focus in pathologic gait analysis (the fact that the patient has CP is known

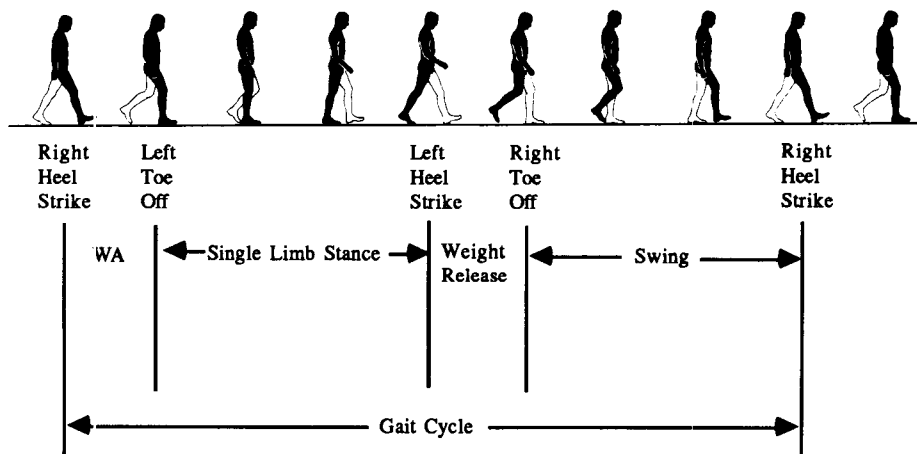


FIGURE 1. Important events and phases in the gait cycle. (Modified from Inman et al., 1981, p. 26; used with permission of the publisher.)

before gait analysis is performed). These effects include muscle tightness, spasticity, and weakness, all of which affect the patient's gait motions.

Hence, the goal of diagnosis in this domain is to identify the improper muscle activity and joint limitations that cause the deviations observed in a patient's gait. The input is the information gathered by a gait analysis laboratory. There are three types of data: clinical, historical, and motion. Clinical data result from physical examination of the patient and determine both the range of motion of the different joints and a qualitative measure of the strength of different muscle groups. Historical data include information about any past medical procedures or diagnoses. Motion data specify the time/distance parameters of walking (velocity, stride length, stance and swing times, etc.) and the angular position of the patient's joints (hips, knees, and ankles) during the different gait phases, with the latter recorded in all three planes. Motion data also include electromyograph (EMG) data on selected muscle groups, indicating when nervous stimulation occurs during the gait cycle. Figure 2 shows what some of this information look like. Typically, these data are gathered before a gait analysis is performed.


This domain is complex for a number of reasons. Patients with neurological disorders such as cerebral palsy have a wide variation of muscle and joint faults, and typically each patient has multiple faults. Reasoning about multiple faults is difficult because gait involves a number of highly interacting components and processes. The domain is further complicated because the system attempts to compensate for faults. Furthermore, many gait parameters cannot be directly measured with current technology. For example, EMG data are at best a qualitative measure of muscle forces (Simon, 1982).

Historical Data

(both hamstrings lengthened 2/88. Mother reported post-operative improvement of gait, but the hamstrings are tightening again.)

	R	L	Comments
Hip			
flex	WN / 65	WN / 75	
ext	-15	-15	← -30 mm
abd			
add			
I/E rota.	/	/	
Knee			
flexion	/	/	
extension			
Ankle			
dorsiflex	10 / 5	10 / 5	
plantarflex			
in/eversion	/	/	
Forefoot			
add/abd	/	/	
pro/sup	/	/	
Toes			
flex/ext			

Trunk *varus tendency of ankles*



Motion Data

Clinical Data (Range of motion measurement)



EMG Data

FIGURE 2. Some of the data used in gait analysis.

PREVIOUS GAIT ANALYSIS PROGRAMS

Most of the medical research on pathological gait is either statistical in nature (Wong et al., 1983) or concentrates on the functionality of a particular joint or muscle group (Simon et al., 1978). There has not been a good attempt to formalize a method for gait analysis. Most of the bioengineering research on formalizing gait is based on quantitative models (Hemami, 1985). These models help one to understand gait, but they do not help analyze the data. However, there have been two notable attempts to create gait analysis programs using artificial intelligence methods.

In the mid-1970s, Tracy and others at Stanford University developed an expert system to help diagnose gait disorders (Tracy et al., 1979). This program relies on associational knowledge encoded as MYCIN-like production rules (Shortliffe, 1976). Using the MYCIN framework, it was difficult to direct the

reasoning in an organized manner, that is, across phases or joints. Below, we shall see that DR. GAIT-1 and its successors carefully control the sequence of reasoning. In addition, Tracy's program considers only muscle tightnesses and weaknesses, ignoring other possible causes such as muscle spasticities and joint contractures. Finally, it does not relate the faults to the original deviations.

In the early 1980s at Vanderbilt University, Dzierzanowski (1984) and others built an expert system called GAITSPERT. GAITSPERT attempts to categorize patients according to different patterns of motion. It first attempts to identify general patterns of motion, and then, based on which general patterns match, more specific patterns are examined. GAITSPERT controls its reasoning better than Tracy's program but is limited to the domain of stroke-related disorders. By performing gait analysis for cerebral palsy patients, DR. GAIT-1 and DR. GAIT-2 not only add more breadth to the coverage of gait analysis problems but also deal with more complex problems. As we shall see, DR. GAIT-2 explicitly reasons about multiple, interacting faults. Our third system, QUAWDS moves away from the limitation to a particular neurological or orthopedic disorder.

DR. GAIT-1

DR. GAIT-1 is the first of the expert systems in pathologic gait analysis that we have developed. This system relies exclusively on associational knowledge; no explicit causal reasoning is performed. DR. GAIT-1 operates strictly by associating patterns of observations with causes.

Functional Organization of DR. GAIT-1

DR. GAIT-1 analyzes the motion of one leg in one plane, specifically the angular positions of the hip, knee, and ankle in the sagittal plane, which is the view from the side. The primary inputs are scaled motion data and interpreted EMG data. The motion data are grouped by the four phases of gait with single limb stance and swing split into two parts and is scaled by 5° increments and decrements from normal. The scale ranges from markedly decreased ($25^\circ <$ normal) to markedly increased ($25^\circ >$ normal). Figure 3 illustrates scaling of motion. The scaling was performed by hand for DR. GAIT-1 but was automated when DR. GAIT-2 was implemented.

The EMG data are interpreted to determine if a given muscle was on or off in a particular phase. Figure 4 shows an example EMG interpretation. This interpretation is performed by a domain expert. Automation of EMG interpretation is being investigated.

Based on these data, DR. GAIT-1 does diagnosis by performing three sub-tasks:

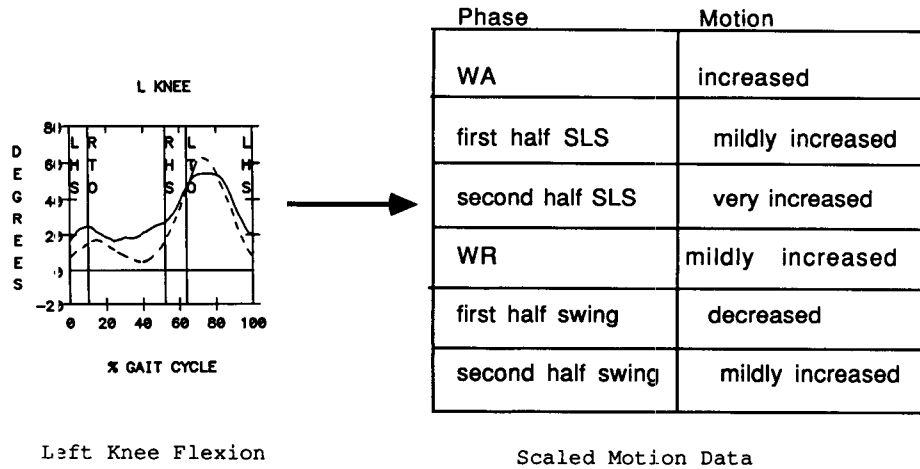


FIGURE 3. Example scaling of motion.

1. Match patterns. For each joint, match the pattern of motion across the phases to a set of precompiled patterns of motion.
2. Match faults. Using the motion patterns plus information about EMG, determine the general faults of the patient.
3. Specialize faults. The descriptions of the faults are specialized to correspond to the observed motions and EMG activity of the patient.

Each of these tasks is performed by a set of rules that directly maps inputs to outputs. These rules were implemented using the GENIE knowledge engineering

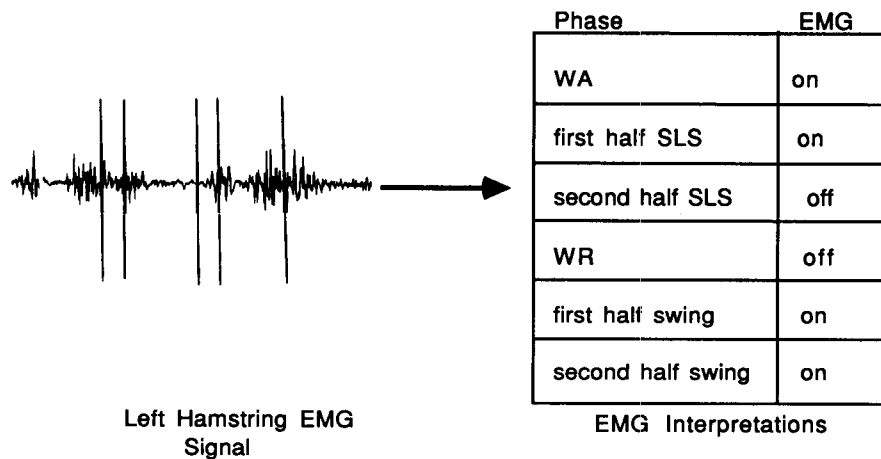


FIGURE 4. Example interpretation of EMG.



FIGURE 5. Functions performed by DR. GAIT-1.

tool, which combines frames and rule bases (Sandell, 1984). Figure 5 illustrates these functions.

Figure 6 illustrates a rule for each of DR. GAIT-1's diagnosis subtasks. They are stated in English for the convenience of the reader. The first rule looks for a particular pattern of ankle motion, namely whether the ankle has increased dorsiflexion (the foot directed toward the shank more than normal) during SLS. Based on this pattern, the second rule will conclude that the gastroc/soleus muscle (the calf muscle) is weak. Note that the rule does not require an exact match of the pattern, since other conditions of the rule consider specific motions. The third rule specializes this diagnosis if increased dorsiflexion occurs only during the first half of single limb stance. If this chain of rules fires, DR. GAIT-1 will reach the following conclusion:

Abnormal dorsiflexion during the first half of single limb stance is noted.
Gastroc/soleus activity is unable to counteract body weight dorsiflexion torque.

Analysis of DR. GAIT-1

DR. GAIT-1 was informally tested on 20 cases. The testing included comparing the system's performance with that of a domain expert over the tasks of identify-

- if the ankle position during WA is within normal range, and
the ankle position during the first half of SLS is at least mildly increased, and
the ankle position during the second half of SLS is at least mildly increased, and
the ankle position during WR is within normal range, and
the ankle position during the first half of swing is within normal range, and
the ankle position during the second half of swing is within normal range;
then conclude pattern of abnormal dorsiflexion during SLS.
- if there is a pattern of abnormal dorsiflexion during SLS, or
the ankle position during the first half of SLS is at least mildly increased, or
the ankle position during the second half of SLS is at least mildly increased;
then conclude weak gastroc/soleus muscle causing abnormal dorsiflexion during SLS.
- if there is a weak gastroc/soleus muscle causing abnormal dorsiflexion during SLS, and
the ankle position during the first half of SLS is at least mildly increased, and
the ankle position during the second half of SLS is not at least mildly increased;
then specialize diagnosis to first half of SLS.

FIGURE 6. Example rules of DR. GAIT-1.

ing motion deviations and then identifying the deviations' causes. On simple cases—60% of the test cases—DR. GAIT-1 identified 80% of the major deviations and identified the causes correctly.

However, DR. GAIT-1 has difficulties with harder cases, for several reasons. First, DR. GAIT-1 uses only empirical pattern matching in its problem solving. Only patients whose symptoms match exactly the situations described by the rules can have their gait adequately analyzed by the program. Adding new rules to cover each new specific situation is not an adequate solution because there is a combinatorial number of multiple fault possibilities.

A second problem is that nothing checks the consistency of the program's conclusions with the patient data. If a rule concludes hamstring overactivity, this is given as an answer regardless of whether additional data would discredit this hypothesis.

A related problem is that the explanations offered by the rule base are poor. The only types of explanations the system can give are run-time trace explanations. These explanations detail the sequence of diagnostic reasoning of the case and identify how certain observations or problem states match the knowledge base. However, the system is unable to justify its conclusions in terms of what motions are caused by which faults.

Finally, another problem is lack of intermediary concepts within the system. The system is always matching a set of observations directly to a fault. DR. GAIT-1's rules do not express concepts common to a large number of situations.

To overcome these problems, an underlying domain model is needed that can determine interactions in the domain and formulate reasonable explanations. The key to doing this is to use some understanding about how gait is caused, namely that the joints' motions are caused by a combination of torques produced by muscles, body weight, and momentum.

DR. GAIT-2

If a gait analysis system could reason about the combined effects of muscles, joints, weight, and momentum on joint rotation, it would be able to propose and evaluate faults based on a causal understanding of the domain. The opportunity then exists to focus on particular abnormal motions and consider only fault hypotheses that are causally relevant. The opportunity also exists to determine the *explanatory coverage* of a fault hypothesis, that is, whether it accounts for the observed gait. Observations that are not accounted for can become the focus for further reasoning. Because of its capability for causal reasoning, DR. GAIT-2 is a better, more robust system than DR. GAIT-1.

Functional Organization of DR. GAIT-2

Just like DR. GAIT-1, DR. GAIT-2 analyzes only the hip, knee, and ankle motions of one leg in the sagittal plane (side view). Again as with DR. GAIT-1, the primary inputs are history, physical exam, scaled motion data, and interpreted EMG data. To do diagnosis, DR. GAIT-2 does the following series of subtasks:

1. Identify deviations. A motion deviation must be 10° or more from normal to be important enough to explain. As in DR. GAIT-1, these are grouped by joint (hip, knee, and ankle) and phase, with SLS and swing again split into two parts.

2. Diagnose classes of causes. DR. GAIT-2 determines the classes of causes that exist before considering specific faults. There are three classes of causes:

- (a) Limited range of motion. This is associated with very restricted ranges of motion throughout the gait cycle by any of the joints, such that the motion can be attributed to co-contractions of opposing muscles. To conclude this class, the patient data must indicate that the opposing muscles are continuously active.

- (b) Contracture. This is associated with restricted motion of a joint throughout the gait cycle caused by tight contracted muscles or tight joint capsules. To conclude this class, the contracture must be specified in the clinical patient data.

- (c) Dynamic. If an abnormal motion is not explained by either of the above two classes, its cause is considered to be dynamic, that is, caused by the particular dynamics of muscle actions, body weight, and momentum during that phase of the gait.

3. Diagnose dynamic causes. DR. GAIT-2 uses its causal model of gait to generate and select hypotheses. An assumption-based truth maintenance system (de Kleer, 1986) is used to ensure that no hypotheses conflict with each other.

These functions were implemented using the GENIE tool. Figure 7 illustrates these functions. The operation of the causal model is described in the following two sections.

Causal Reasoning in DR. GAIT-2

DR. GAIT-2's causal reasoning about torques is the heart of the system. As mentioned above, the rotational motion at a joint is the result of the combination of torques acting on the joint. The knee's motion, for example, is determined by all the torques acting on the knee. For the knee, DR. GAIT-2 reasons about the




FIGURE 7. Functions performed by DR. GAIT-2.

torques caused by the hamstrings (muscles on the back of the thigh), quadriceps (muscles on the front of the thigh), gastroc/soleus (the calf muscles), body weight, and body momentum. If the knee's position is abnormal, DR. GAIT-2 infers that at least one of the torques is abnormal.* For example, if the knee shows increased flexion, one possibility is an increased hamstring torque.

The total torque acting on the knee, then, must satisfy this equation:

$$\text{knee-torque} = \text{hamstring-torque} + \text{quadricep-torque} + \text{gastroc/soleus-torque} + \text{bodyweight-torque} + \text{momentum-torque} \quad (1)$$

However, now five torques must be considered to assign the blame for abnormal knee position. Unfortunately, the equation cannot be straightforwardly solved because numeric measurements of the various torques are not available. Consequently, qualitative reasoning is called for, but the equation as it stands is underconstrained—an increase in any one or any combination of the torques could account for increased knee flexion. To resolve these problems, DR. GAIT-2 uses case data, general knowledge about gait, and heuristics about which abnormalities are more likely.

In most problems involving search and combination, it helps to organize the search space. Each torque on the knee can be classified as a *flexion torque* or an *extension torque* based on whether the torque normally causes flexion or extension, respectively. For the knee, the quadricep-torque is a flexion torque and the hamstring-torque and gastroc/soleus-torque are extension torques. The classification for bodyweight-torque and momentum-torque depends on the phase of the gait. We also classify a torque as *internal* if it is produced locally or *external* otherwise. All muscle torques are internal, while the bodyweight and momentum torques are external. Using these categories, the torques can be organized as a tree as shown in Fig. 8.[†]

The tree in Figure 8 organizes Equation (1) as the following set of equations:

$$\text{knee-torque} = \text{flexion-torque} - \text{extension-torque} \quad (2)$$

*In an accurate physical model, abnormal torque results in abnormal angular acceleration, which usually, but not always, results in abnormal angular position. One of the goals of the QUAWDS system is to reason about the intermediate concept of angular acceleration.

[†]The "bodyweight-torque" in the figure includes the effects of both body weight and momentum.

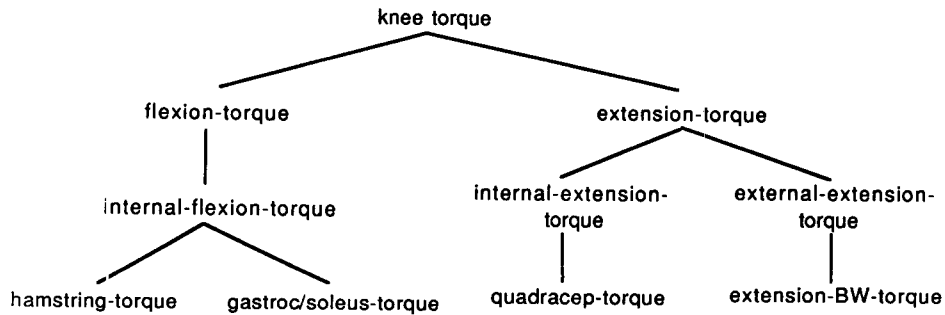


FIGURE 8. Torque tree for knee in second half of swing.

$$\text{flexion-torque} = \text{internal-flexion-torque} \tag{3}$$

$$\text{extension-torque} = \text{internal-extension-torque} + \text{external-extension-torque} \tag{4}$$

$$\text{internal-flexion-torque} = \text{hamstring-torque} + \text{gastroc/soleus-torque} \tag{5}$$

$$\text{internal-extension-torque} = \text{quadricep-torque} \tag{6}$$

$$\text{external-extension-torque} = \text{extension-BW-torque} \tag{7}$$

Now that DR. GAIT-2 has all of these equations, how can they be used? The scaled gait motions indicate whether a joint's position is increased, decreased, or normal. Similarly, the torques are described as increased, decreased, or normal. DR. GAIT-2 solves the equations using these qualitative values by employing de Kleer's incremental qualitative (IQ) algebra (de Kleer, 1979). The rules for IQ addition are shown in Table 1. In the table, ↑ stands for increased, 0 stands for normal, ↓ stands for decreased, and ? stands for unknown.

In the two unknown cases in Table 1, more information is needed to disambiguate the answer. To do this, we introduce the relations <diff, >diff, and =diff. The statement $A < \text{diff } B$ says that "A has a smaller deviation from normal than B." Thus, if A is increased, B is decreased, and $A < \text{diff } B$, then $A + B$ is decreased; that is, B's decrease from normal is greater than A's

TABLE 1. IQ Addition Table

+	↓	0	↑
↓	↓	↓	?
0	↓	0	↑
↑	?	↑	↑

increase from normal. Using these relations, we have constructed a modified IQ addition table, which is shown in Table 2.

Since no numeric measurements for the torques are available, it is uncertain whether a torque is increased or decreased and whether the relationship between two torques is $<$ diff, $=$ diff, or $>$ diff. However, hypotheses about the amount of and relationships between torques can be made. In the context of a single torque tree, DR. GAIT-2's strategy is to hypothesize everything that is physically reasonable and then use heuristic knowledge about cerebral palsy to choose the best set of hypotheses. This hypothesis set corresponds to the diagnosis for that joint and phase. Since there are other torque trees for other joints and phases, DR. GAIT-2 must in addition ensure that all of the hypotheses that are chosen are consistent with each other. To maintain consistency, a support system based on de Kleer's (1986) assumption-based truth maintenance system (ATMS) is used. That is, each set of hypotheses produced by processing a torque tree corresponds to a set of assumptions. Thus, by maintaining consistency between assumptions, the ATMS also maintains consistency of the overall diagnosis.

Example of Causal Reasoning in DR. GAIT-2

How does DR. GAIT-2 apply a torque tree to patient data? The following example demonstrates how this is done.

Suppose that during the second half of swing the patient's knee has increased flexion, and the following data describe muscle activity:

Muscle	Usual activity	Actual activity
Gastroc/soleus	Off	On
Hamstrings	On	On
Quadriceps	Off	Off

"Usual activity" indicates normal muscle activity, while "Actual activity" is the EMG interpretation for the patient. Figure 8 is the relevant torque tree in this situation and Eqs. (2)–(7) are the relevant torque equations.

First, DR. GAIT-2 uses domain knowledge to determine the possible values of the lowest level torques (muscle torques and body weight torque), which are as follows:

$$\text{hamstring-torque} \in \{ \langle \text{increased}, \{ \text{hamstring not weak} \} \rangle, \langle \text{decreased}, \{ \text{hamstring weak} \} \rangle \}$$

$$\text{gastroc/soleus-torque} \in \{ \langle \text{increased}, \{ \} \rangle \}$$

TABLE 2. Modified IQ Addition Table

		<i>B</i>			
		+	↓	0	↑
<i>A</i>	↓	↓	↓	↓	case2
	0	0	↓	0	↑
	↑	↑	case1	↑	↑

$$\text{case1} = \begin{cases} \uparrow & \text{if } A > \text{diff } B \\ 0 & \text{if } A = \text{diff } B \\ \downarrow & \text{if } A < \text{diff } B \end{cases} \quad \text{case2} = \begin{cases} \uparrow & \text{if } A < \text{diff } B \\ 0 & \text{if } A = \text{diff } B \\ \downarrow & \text{if } A > \text{diff } B \end{cases}$$

$$\text{quadricep-torque} \in \{ \langle \text{normal}, \{ \} \rangle \}$$

$$\text{extension-BW-torque} \in \{ \langle \text{decreased}, \{ \} \rangle \}$$

Each torque value has the form $\langle V, \{A_1, A_2, \dots\} \rangle$ where V is increased, normal, or decreased and each A_i is a hypothesis about a torque or torques; each A_i is treated as an assumption by the ATMS. At this level, the hamstring-torque could conceivably be increased or decreased because the hamstrings could either be overly weak or strong. The gastroc/soleus-torque is inferred to be increased because the gastroc/soleus is on when it is normal to be off. The quadricep-torque is inferred to be normal because the quadriceps is off as it is normally should be. The extension-BW-torque is inferred to be decreased based on the conclusions of heuristic rules which use other patient data not described here.

These torque values are then propagated up to the joint's torque. At the next level of the torque tree in Fig. 8, the following values are produced:

$$\begin{aligned} \text{internal-flexion-torque} \in & \{ \langle \text{increased}, \{ \text{hamstring not weak} \} \rangle , \\ & \langle \text{increased}, \{ \text{gastroc/soleus-torque} > \text{diff hamstring-torque}, \\ & \text{hamstring weak} \} \rangle , \\ & \langle \text{normal}, \{ \text{gastroc/soleus-torque} = \text{diff hamstring-torque}, \\ & \text{hamstring weak} \} \rangle \} \\ & \langle \text{decreased}, \{ \text{gastroc/soleus-torque} < \text{diff hamstring-torque}, \\ & \text{hamstring weak} \} \rangle \} \\ \text{internal-extension-torque} \in & \{ \langle \text{normal}, \{ \} \rangle \} \\ \text{external-extension-torque} \in & \{ \langle \text{decreased}, \{ \} \rangle \} \end{aligned}$$

Note that if the hamstrings are hypothesized to be weak, it is possible that the increase in gastroc/soleus-torque could exceed, equal, or fall short of the decrease in hamstring-torque. These possibilities are generated using the modified IQ addition table in Table 2.

At the next to last level, flexion-torque has exactly the same possibilities as internal-flexion-torque. Extension-torque is decreased because external-extension-torque is decreased.

Finally, DR. GAIT-2 generates the possible values for knee-torque:

$$\text{knee-torque} \in \{ \langle \text{increased, } \{\text{hamstring not weak}\} \rangle, \\ \langle \text{increased, } \{\text{gastroc/soleus-torque} > \text{diff hamstring-torque,} \\ \text{hamstring weak}\} \rangle, \\ \langle \text{increased, } \{\text{gastroc/soleus-torque} = \text{diff hamstring-torque,} \\ \text{hamstring weak}\} \rangle, \\ \langle \text{increased, } \{\text{flexion-torque} < \text{diff extension-torque,} \\ \text{gastroc/soleus-torque} < \text{diff hamstring-torque,} \\ \text{hamstring weak}\} \rangle, \\ \langle \text{normal, } \{\text{flexion-torque} = \text{diff extension-torque,} \\ \text{gastroc/soleus-torque} < \text{diff hamstring-torque,} \\ \text{hamstring weak}\} \rangle, \\ \langle \text{decreased, } \{\text{flexion-torque} > \text{diff extension-torque,} \\ \text{gastroc/soleus-torque} < \text{diff hamstring-torque,} \\ \text{hamstring weak}\} \rangle \}$$

Note that because Eq. (2) has a minus sign instead of a plus sign, the decrease in extension-torque corresponds to an increase to knee-torque. Thus, all the possible values of flexion-torque that are increased or normal are possible ways for knee-torque to be increased. If flexion-torque is decreased, again Table 2 must be used to generate possible values of knee-torque based on the possible relationships between flexion-torque and extension-torque.

Imposing the known constraint that the knee-torque is increased causes the system to remove inconsistent values. This leaves four possible sets of hypotheses that account for increased knee flexion in the second half of swing. To select one of the sets, each set is scored by using domain heuristics that score each hypothesis on the basis of how likely it is to occur in cerebral palsy patients. For example, for CP patients it is unlikely that a muscle is weak (unless surgery has been performed on the muscle), so hypotheses about muscle weakness will receive a high score (higher means less likely). On the other hand, it is very likely that a muscle is overactive, so these hypotheses will receive low scores. The score for a set of hypotheses is the sum of the scores of its elements. The lowest scoring set (the most likely one) is selected to be the best possible explanation of the torque's value. Based on these factors, the top scoring set of hypotheses in the example is {hamstring not weak}, and DR. GAIT-2 provides the following diagnosis and explanation of the abnormal motion:

Problem name: right-knee-sagittal-second-half-swing-flexion

Problem summary:

The right knee has increased flexion during second-half-swing.

Assuming the following:

1) (patient data right hamstring muscle-strength) equals nonweak.

The PRIMARY CAUSE(s) of this problem is(are):

increased hamstring-torque. which is due to normal-firing
of a functionally-spastic hamstring.

increased gastroc-soleus-torque. which is due to abnormal-firing
of a functionally-spastic gastroc/soleus.

The AUXILIARY CAUSE(s) of this problem is(are):

decreased extension-BW-torque

The phrase “increased hamstring-torque . . . due to normal-firing . . .” indicates that it is normal for the hamstrings to be firing but that the hamstrings are producing more torque than normal. The hamstring is said to be spastic because overactive muscles imply spasticity in cerebral palsy patients. Note that DR. GAIT-2 remembers that the gastroc/soleus was on when it normally should have been off and also that the torque caused by body weight and momentum is different from normal.

It is possible that “hamstring weak” will be selected to diagnose some other abnormal motion of the patient. In this situation, the ATMS will discover the contradiction, which results in DR. GAIT-2 constructing two alternative diagnoses. For each abnormal motion in which “hamstring not weak” was selected, DR. GAIT-2 constructs a new diagnosis by selecting the best sets of hypotheses that do not make this assumption. Constructing the other diagnosis is similar, except that DR. GAIT-2 selects sets of hypotheses that do not include “hamstring weak.” The two alternative diagnoses are compared via their scores and the best one is selected.*

Analysis of DR. GAIT-2

We tested DR. GAIT-2 on 22 cases covering a range of cerebral palsy patients. The overall set of cases was more difficult than the set of cases used to test DR. GAIT-1. The program’s performance at identifying abnormal motions and explaining their causes was compared with the written reports generated by the domain expert.

In the 22 cases, 170 abnormal motions were mentioned in the reports. DR.

*It is possible that only one of the alternatives can be constructed; for instance, “hamstring not weak” might be necessary to account for some abnormal motion. It is also possible that no alternative can be constructed or that the alternatives contain contradictions among other assumptions. In these cases, DR. GAIT-2 is unable to continue.

GAIT-2 identified 89% (151) of these abnormal motions. Most of the omissions are range-of-motion problems, apparently because the triggering conditions for this class of abnormal motions are too restrictive. DR. GAIT-2 also identified 46 abnormal motions not mentioned in the reports. Most of these additional problems are minor or were perceived to be insignificant.

At identifying the causes of finding, the system found the correct causes 95% of the time (it was correct for 187 of the 196 abnormal motions it found). Most of the mistakes occurred because DR. GAIT-2 doesn't know to what degree particular muscles can influence the various joints. The other errors resulted from incorrect modeling of body weight at the knee during WA.

It appears, then, that DR. GAIT-2 is very successful at identifying abnormal motions and diagnosing their causes. With some refinements to the knowledge base, it is possible that its performance on these tasks could be even better. The improved performance over DR. GAIT-1 can be directly attributed to the causal model of the domain. DR. GAIT-2 is able to overcome many of the holes in DR. GAIT-1's knowledge by deriving the relationships between observations and faults rather than relying solely on precompiled associations.

Nevertheless, DR. GAIT-2 still has several limitations. The representation of time in DR. GAIT-2 is very elementary. The gait cycle is divided into a fixed number of phases and each phase is treated as a single point of time. This temporal representation makes it hard to specify intervals of interest by the actions and events of a particular patient's gait.

The causal model does not consider several factors that determine the relative amount of torque that a muscle can produce. For example, the torque of a muscle is affected by the joint's position. Also, the model does not recognize the relative strengths of opposing muscles acting on a joint.

If DR. GAIT-2 determines that there is a joint contracture, the causal model is not applied to that joint. The contracture should be represented as a special kind of torque that occurs only when the joint's position is at the limit of its range of motion.

Furthermore, DR. GAIT-2 is limited in its domain: analyzing the motions of one leg in the sagittal plane in a single visit by a patient with cerebral palsy. Human gait involves coordination between both legs, and although sagittal plane motion is the most important, movements in other planes affect one's gait. Patients are often analyzed more than once, such as before and after treatment; it would be useful to determine how the patient's gait has (or has not) improved. Also, other types of disorders affect gait, including stroke, head injuries, arthritis, muscular dystrophy, and fractures with subsequent complications.

Finally, DR. GAIT-2's reasoning is inefficient in some respects. When it considers a deviation, it always generates all possible sets of fault hypotheses that could account for the deviation. DR. GAIT-2 does not check if fault hypotheses that have already been selected could account for the deviation. Also, it does

not try to generate more likely hypotheses first, but considers everything that is remotely plausible.

QUAWDS

We are developing another knowledge-based system, called QUAWDS, to resolve many of the limitations of DR. GAIT-2. In particular, QUAWDS is intended to perform gait analysis for both legs, all three planes, and several neurological and orthopedic disorders that affect gait. It would be difficult to extend DR. GAIT-2 to achieve this goal. While causal reasoning is a necessity for diagnosis of pathologic gait, causal reasoning without any guidance from associational knowledge is computationally complex (Bylander et al., 1989). Some of these elements already exist in DR. GAIT-2. In particular, DR. GAIT-2 uses associational knowledge to select among the hypotheses the torque trees generate. However, this knowledge is not organized with the purpose of controlling causal reasoning.

Functional Organization of QUAWDS

The organization of QUAWDS will resemble DR. GAIT-2 in several ways. QUAWDS also has to identify deviations and diagnose faults that account for the deviations. However, it is organized so that associational knowledge provides information to control causal reasoning. QUAWDS will perform the following subtasks in doing diagnosis:

1. Identify deviations. Like DR. GAIT-2, a motion deviation must be 10° or more from normal to be important enough to explain. Unlike DR. GAIT-1 or DR. GAIT-2, the duration of a deviation is not limited to a small number of phases. QUAWDS will be able to describe a deviation as occurring over several phases or a fraction of a phase.
2. Generate causes of deviations. Like DR. GAIT-2, QUAWDS will, for a specific deviation, use its qualitative causal model to generate fault hypotheses that can account for the deviation.
3. Score faults. Each fault will be associated with a collection of heuristic rules that can score the faults. As noted above, DR. GAIT-2 also did this, but QUAWDS will also know about classes of faults corresponding to the legs, joints, and muscles. Another new feature is that these rules will explicitly incorporate conditions about what disorder (e.g., cerebral palsy or stroke) is the original cause of the faults. This will give the reasoning additional focus. For example, if the patient has left hemiplegia (injury to the right side of the brain), faults of the left leg should be considered before those of the right leg.
4. Determine explanatory coverage. Assuming a particular fault or set of




FIGURE 9. Functions performed by QAWDS.

faults, QAWDS' causal model will determine the deviations that can and cannot be accounted for.

5. Assemble a diagnosis. The above subtasks will not be run in a simple sequence but will be controlled by an *abductive assembler* (Josephson et al., 1987). The assembler iteratively performs the following steps: select the most significant deviation not yet accounted for, find out what faults can account for it, select the fault with the best score, determine what deviations can be accounted for by the faults that have been selected, and repeat until as much is accounted for as possible. The assembler's procedure is similar to how DR. GAIT-2 operates but focuses attention on the most important finding not yet accounted for.

We will be using generic tasks, a set of tools for modularizing knowledge-based systems, to implement these functions (Chandrasekaran, 1986). Figure 9 illustrates the organization of these functions. We briefly discuss some of the improvements to the causal model that we are developing.

Improvements in the Causal Model

QAWDS' qualitative causal model will start with the same torque equations as DR. GAIT-2, except that the equations will have additional elements.* Two elements not included in DR. GAIT-2's equations include range-of-motion constraints and interactions between joints. Whenever a joint reaches a limit on its range of motion, it generates an opposing torque. A joint with limited range of motion, that is, a joint contracture, is a special case of this effect. Also, motions at one joint will affect nearby joints. For example, ankle plantar flexion during stance will contribute to knee extension.

*Also, there will be equations for rotational motions not represented by Dr. GAIT-2.

QUAWDS will use the causal model differently by directly relating the equations to angular acceleration, as opposed to angular position. Given a deviation (a position different from normal), QUAWDS will trace this to abnormal angular acceleration previous to the deviation. Fault hypotheses will be generated based on the abnormal acceleration rather than the abnormal position.

In addition, QUAWDS will use the causal model to evaluate its working diagnosis as hypotheses are added to it. The torque equations can be used to check whether a set of faults is consistent with the patient's motion. If it is not, QUAWDS will determine what motions are not accounted for, which can then be used to search for additional fault hypotheses. Thus, QUAWDS will never generate hypotheses based on motions that have already been accounted for.

Of course, because QUAWDS is still in development, we should caution the reader that the above improvements are promises yet to be fulfilled. Nevertheless, by building on the success of DR. GAIT-2, we believe that QUAWDS will represent the next step for gait analysis programs.

CONCLUSION

We have presented our work on three diagnostic expert systems in the domain of gait analysis: DR. GAIT-1, DR. GAIT-2, and QUAWDS. The success and the limitations of the first two systems have led to the development of their successors.

DR. GAIT-1 relies solely on associational knowledge and so is able to succeed only on situations that it was able to match exactly. DR. GAIT-1 shows that diagnosis of pathologic gait is a complex task that requires causal reasoning.

With its qualitative causal model, DR. GAIT-2 succeeds in a limited subdomain of gait analysis. It shows that causal models can be effectively used to produce high-quality diagnoses in complex domains. By representing and reasoning about torques, DR. GAIT-2 is able to construct diagnoses that are causally relevant to the patient's abnormal motions. An ATMS ensures that the diagnosis contained no contradictions. However, much of DR. GAIT-2's success is also because associational knowledge is able to make choices among causally equivalent alternatives.

QUAWDS, the system we are currently developing, will apply to a larger subdomain of gait analysis. QUAWDS more clearly identifies the role of associational knowledge in causal reasoning and will make several improvements to DR. GAIT-2's causal model.

No doubt QUAWDS will have major limitations as a gait analysis program. Nevertheless, we are confident that we will continue to make incremental progress in developing gait analysis programs and, more generally, in understanding how to use causal reasoning to perform diagnosis in complex domains.

REFERENCES

- Bylander, T., Allenmang, D., Tanner, M. C., and Josephson, J. R. 1989. Some Results Concerning the Computational Complexity of Abduction. Proc. First Int. Conf. on Principles of Knowledge Representation and Reasoning, Toronto, pp. 44-54.
- Bylander, T., Weintraub, M. A., Simon, S. R. 1988. Design of an Expert System for Gait Analysis. Technical report, Laboratory for AI Research, CIS Dept., Ohio State Univ., Columbus.
- Chandrasekaran, B. 1986. Generic tasks in knowledge-based reasoning: High-level building blocks for expert system design. *IEEE Expert* 1(3):23-30.
- de Kleer, J. 1979. The origin and resolution of ambiguities in causal arguments. *Proc. Sixth Int. Joint Conf. on Artificial Intelligence*, Tokyo, pp. 197-203.
- de Kleer, J. 1986. An assumption-based TMS. *Artif Intell* 28(2):127-162.
- Dzierzanowski, J. 1984. Artificial Intelligence Methods in Human Locomotor Electromyography, Ph.D. thesis, Vanderbilt University.
- Hemari, H. 1985. Modeling, control, and simulation of human movement. *Crit Rev Biomed Eng* 13(1):1-34.
- Hirsch, D. E. 1987. An Expert System for Diagnosing Gait for Cerebral Palsy Patients. Technical Report MIT/LCS/TR-388, Laboratory for Computer Science, MIT, Cambridge, Mass.
- Inman V. T., Ralston, H. J., and Todd, F. 1981. *Human Walking*. Williams & Wilkins, Baltimore.
- Josephson, J. R., Chandrasekaran, B., Smith, J. W., and Tanner, M. C. 1987. A mechanism for forming composite explanatory hypotheses. *IEEE Trans Syst Man Cybernet* 17(3):445-454.
- Perry, J. 1985. Normal and pathologic gait. In *Atlas of Orthotics*, chapter 4. St. Louis: Mosby.
- Sandell, H. 1984. GENIE User's Guide and Reference Manual. Technical Report 84-003, Electrical and Biomedical Engineering, Vanderbilt University.
- Shortliffe, E. H. 1976. *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier.
- Simon, S. R. 1982. Kinesiology—Its measurement and importance to rehabilitation. In V. L. Nickel, editor, *Orthopedic Rehabilitation*, chapter 5. New York: Churchill Livingstone.
- Simon, S. R., Mann, R. A., Hagy, J. L., and Larson, L. J. 1978. Role of the posterior calf muscles in normal gait. *J Bone Joint Surg* 60:465-472.
- Tracy, K., Montague, E., Gabriel, R., and Kent, B. 1979. Computer-assisted diagnosis of orthopedic gait disorders. *Phys Therapy* 59(3):268-277.
- Wong, M., Simon, S. R., and Olshen, R. 1983. Statistical analysis of gait patterns of persons with cerebral palsy. *Stat Med* 2:345-354.

Request reprints from Tom Bylander.