Using Causal Reasoning in Gait Analysis[1]

David E. Hirsch*
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology

Sheldon R. Simon
Department of Surgery
The Ohio State University

Tom Bylander and Michael A. Weintraub
Department of Computer and Information Science
The Ohio State University

Peter Szolovits
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology

This is a draft printed May 12, 1995.

Abstract

An outstanding abstract will be put here.

*David E. Hirsch is now at Price Waterhouse.

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

***We acknowledge everybody and anybody that we should.

[1]This is a nearly-final draft of the paper that appeared as D. E. Hirsch, S. R. Simon, T. Bylander, M. A. Weintraub, and P. Szolovits. Using causal reasoning in gait analysis. *Applied Artificial Intelligence*, 3(2-3):337-356, 1989. The final version is not available on-line.

Using Causal Reasoning in Gait Analysis

David E. Hirsch, Sheldon R. Simon,
Tom Bylander, Michael A. Weintraub, and Peter Szolovits

1 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 provided by one of the authors (Simon). The difference between each successive system is the recognition of the need for both causal reasoning about the process of gait and for experiential reasoning that can control the complexity of causal reasoning.

The first system (Dr. Gait-1), developed by three of us (Hirsch, Simon, and Szolovits) at MIT and Harvard, relies exclusively on experiential domain models. Dr. Gait-1's success in diagnosing cases is quite limited. The second system (Dr. Gait-2), also 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 computational complexity. The third system (Quawds, for QUalitative Analysis of Walking Disorders), currently being developed by three of us (Bylander, Weintraub, and Simon) at The Ohio State University, is an experiment in integrating causal reasoning with experiential reasoning 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. Due to 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 [?]. For further information on Quawds, see Bylander et al. [?].

2 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. 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 phases for each 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 directly affects the brain, and manifests itself by interfering with the control of voluntary motions. The result of cerebral

Figure 1: Important Events and Phases in the Gait Cycle

Figure 2: Some of the Data used in Gait Analysis

palsy on the gait cycle is the improper coordination of muscle activity. It is these effects, and not cerebral palsy itself, that is the focus in pathologic gait analysis (the fact that the patient has CP is known before gait analysis is performed). These effects include muscle tightness, spasticity, and weakness, all of which in turn 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 the 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 angular position of the patient's joints (hips, knees, and ankles) during the different gait phases and are recorded in all three planes. Motion data also includes EMG (electromyograph) data of selected muscle groups, indicating when nervous stimulation occurs during the gait cycle. Figure 2 shows what some of this information looks like. Typically, all of this data is 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

Figure 3: Example Scaling of Motion

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 given current technology. For example, EMG data is at best a qualitative measure of muscle forces [?].

3 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 experiential reasoning, and as such, no explicit causal reasoning is performed. Dr. Gait-1 operates strictly by associating patterns of observations with causes.

3.1 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 is grouped by phases with single limb stance and swing split into two halves 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 is 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. Current research is investigating how to automate EMG interpretation.

Based on this data, Dr. Gait-1 does diagnosis by performing three subtasks:

1. Match Patterns. For each joint, match the pattern of motion across the phases to a set of precompiled patterns of motion.

Figure 4: Example Interpretation of EMG

Figure 5: Functions Performed by DR. GAIT-1

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 for the particular phases in which the faults appear.

Each of these tasks are performed by a set of rules that directly map inputs to outputs. 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 (bent towards 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 only occurs during the first half of single limb stance. If this chain of rules fires, DR. GAIT-1 will reach the following

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

conclusion:

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

3.2 Analysis of Dr. Gait-1

Dr. Gait-1 was informally tested on 20 cases. The testing included comparing the system's performance to a domain expert over the tasks of identifying motion deviations and the identifying the deviation's 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, due to several reasons. First, Dr. Gait-1 only uses 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 are 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, then this is given as an answer regardless of whether there is additional data that would discredit this hypothesis.

A related problem is that the explanations offered by the rule base is 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 to the knowledge base. However, the system is unable to justify its conclusions in terms of how the patient's motions are caused by the hypothesized faults.

Finally, another problem is the lack of intermediary concepts contained within the system. The system is always matching a set of observations directly to a fault. None of Dr. Gait-1's rules capture concepts common across a large number of situations.

Taking all of this into account, an underlying domain model is needed that can determine interactions in the domain and can 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 the combination of several torques.

4 Dr. Gait-2

If a gait analysis system could reason about the combined effects of muscles, joints, weight, and momentum on joint rotation, then 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 those fault hypotheses that are causally relevant. The opportunity also exists to determine the *completeness* of a multiple fault hypothesis, i.e., 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.

4.1 Functional Organization of Dr. Gait-2

Just like Dr. Gait-1, Dr. Gait-2 only analyzes the hip, knee, and ankle motions of one leg in the sagittal plane (side view). Again like Dr. Gait-1, the primary inputs are 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 and phase with SLS and swing split into two halves.

2. Diagnose Classes of Causes. Currently 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-contractures 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 muscles or tight joint capsules. To conclude this class, the contracture must be specified in the patient data.

Figure 7: Functions Performed by DR. GAIT-2

(c) Dynamic. If an abnormal motion is not explained by either of the above two classes, then its cause is considered to be dynamic, i.e., caused by the particular dynamics of the muscles 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 [?] is used to ensure that no hypotheses conflict with each other.

Figure 7 illustrates these functions. The operation of the causal model is described in the following two sections.

4.2 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 upon this joint. The knee's motion, for example, is determined by all the torques acting upon the knee. For the knee, DR. GAIT-2 reasons about the torques caused by the hamstrings (back muscles of the thigh), the quadriceps (front muscles of the thigh), the gastroc/soleus (the calf muscles), body weight, and body momentum. If the knee's position is abnormal, then DR. GAIT-2 infers that at least one of the torques is abnormal.[2] If the knee shows increased flexion, then one possibility is an increased hamstring torque.

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

$$\begin{aligned}\text{knee-torque} = {} & \text{hamstring-torque} + \text{quadricep-torque} + \text{gastroc/soleus-torque} + \\ & \text{bodyweight-torque} + \text{momentum-torque} \qquad (1)\end{aligned}$$

However, now five torques must be considered to assign the blame for abnormal knee position. Unfortunately, the equation can not be straightforwardly solved because numeric

[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.

Figure 8: Torque Tree for Knee in Second Half of Swing

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 Figure 8.[3]

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

$$
\begin{aligned}
\text{knee-torque} &= \text{flexion-torque} - \\
&\quad\ \text{extension-torque} && (2)\\
\text{flexion-torque} &= \text{internal-flexion-torque} && (3)\\
\text{extension-torque} &= \text{internal-extension-torque} + \\
&\quad\ \text{external-extension-torque} && (4)\\
\text{internal-flexion-torque} &= \text{hamstring-torque} + \\
&\quad\ \text{gastroc/soleus-torque} && (5)\\
\text{internal-extension-torque} &= \text{quadricep-torque} && (6)\\
\text{external-extension-torque} &= \text{extension-BW-torque} && (7)
\end{aligned}
$$

Now that Dr. Gait-2 has all of these equations, how can they be used? The scaled gait motions indicates 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

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

+	↓	0	↑
↓	↓	↓	?
0	↓	0	↑
↑	?	↑	↑

Table 1: IQ Addition Table

			B	
	+	↓	0	↑
	↓	↓	↓	case2
A	0	↓	0	↑
	↑	case1	↑	↑

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

Table 2: Modified IQ Addition Table

[?]. 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 <diff B says that "A has a smaller deviation from normal than B". Thus, if A is increased, B is decreased, and A <diff B, then $A + B$ is decreased, i.e., B's decrease from normal is greater than A's 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 assumption-based truth maintenance system (ATMS) was 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.

4.3 Example of Causal Reasoning in Dr. Gait-2

How does Dr. Gait-2 apply a torque tree to patient data? We discuss this question in the context of a particular example. Suppose that during the second half of swing the patient's knee has increased flexion and the following data is 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 Equations 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 bodyweight-torque), which are as follows:

$$\begin{array}{ll} \text{hamstring-torque} & \in \{\ \langle\text{increased, }\{\text{hamstring not weak}\}\rangle, \\ & \quad\ \langle\text{decreased, }\{\text{hamstring weak}\}\rangle\ \} \\ \text{gastroc/soleus-torque} & \in \{\ \langle\text{increased, }\{\}\rangle\ \} \\ \text{quadricep-torque} & \in \{\ \langle\text{normal, }\{\}\rangle\ \} \\ \text{extension-BW-torque} & \in \{\ \langle\text{decreased, }\{\}\rangle\ \} \end{array}$$

Each torque value has the form $\langle V, \{A_1, A_2, \ldots\}\rangle$ 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 increased because the gastroc/soleus is on when it is normal to be off. The quadricep-torque is normal because the quadriceps is off as it is normally should be. The extension-BW-torque is decreased based on 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 Figure 8, the following values are produced:

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

Note that if the hamstrings are hypothesized to be weak, then 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:

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

Note that because Equation 2 uses 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 set, 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, 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
```

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.[4]

4.4 Analysis of DR. GAIT-2

We tested DR. GAIT-2 on 22 cases covering a range of cerebral palsy patients. The program's performance at identifying abnormal motions and explaining their causes was compared to the written reports generated by the domain expert.

In a test set of 22 cases, 170 abnormal motions were mentioned in the reports. DR. 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 out 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 an 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 qualitative torque 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 also recognize the relative strengths of opposing muscles acting on a joint.

If DR. GAIT-2 determines that there is a joint contracture, then the qualitative torque 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, viz., analyzing the motions of one leg in the sagittal plane in a single visit by a patient with cerebral palsy. Human gait involves

[4] It is possible only one of the alternatives can be constructed, e.g., "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.

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, e.g., 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.