

PERFEX: AN EXPERT SYSTEM FOR INTERPRETING 3D MYOCARDIAL PERFUSION

ABSTRACT

Interpreting three-dimensional (3D) data is generally recognized as an ill-defined and information-intensive task. The task becomes increasingly difficult in the context of medical diagnostic imagery, wherein the visual information must be interpreted in conjunction with other, non-visual information. A novel approach is presented to perform the interpretation of such multidimensional information, concentrating on a medically important application: the interpretation of 3D tomograms of myocardial perfusion distribution. The overall goal is to assist in the diagnosis of coronary artery disease. The approach employs knowledge-based methods to process and map the 3D visual information into symbolic representations, which are subsequently used to infer structure (anatomy) from function (physiology), as well as to interpret the temporal effects of perfusion redistribution, and assess the extent and severity of cardiovascular disease both quantitatively and qualitatively. The knowledge-based system presents the resulting diagnostic recommendations in both visual and textual forms in an interactive framework, thereby enhancing overall utility. This article presents the methodology underlying this approach, including the implementation and testing of this system within an actual clinical environment.

Index Terms: Knowledge-Based Systems; medical expert systems; artificial intelligence in medicine.

1. INTRODUCTION

One of the most difficult, important, and ill-defined tasks in medicine is that of interpreting images. It has been pointed out [1] that a radiologist can disagree with his or her own diagnosis (when presented with the same image) as much as 15% of the time. One reason for this is that standardized methodologies for interpreting images generally do not exist, and visual interpretation is thus subject to observer variability. In addition, the decision-making process becomes increasingly difficult since expertise plays a major role, which implies that consistency, reliability, and accuracy are highly dependent on domain knowledge. Furthermore, clinical problem solving typically requires integrating imagery with information from other sources, such as electrocardiographic results, symptoms, and other relevant patient-specific data. Further complexity is added by requiring that the visual information, which commonly appears in distorted or two-dimensional (2D) formats, be visualized and integrated mentally by the clinician into more meaningful forms. Hence, the task of interpreting and visualizing medical imagery can be viewed as an information-intensive process requiring significant medical knowledge and clinical expertise, utilizing diverse types of information that may be distorted, misleading, or possibly inexact.

With this in mind, a knowledge-based (KB) system has been developed to assist in the decision-making process by providing expert-based knowledge with which to process and interpret patient-specific information that is both image- as well as not image-based, while also providing an interactive and clinically useful environment. The approach underlying this system incorporates techniques to accomplish the following: (1) automatically extract features from the 3D imagery, (2) cast the extracted features in symbolic form for subsequent knowledge-based processing; (3) infer anatomical information (i.e., possible locations of arterial lesions) from physiological information (i.e., myocardial perfusion distribution); (4) reason about temporally-derived image information regarding possible perfusion redistribution; (5) integrate other types of non-image, patient-specific information; (6) assess the extent and severity of coronary artery disease; (7) model uncertainty to accommodate inexactness at various levels; (8) assess the overall patient condition; and (9) present all the relevant diagnostic information in a medically meaningful, interactive manner.

This KB system (PERFEX, for perfusion expert) has been under development for several years. The current system design is significantly different from previous ones [2,3] in terms of knowledge base robustness, temporal reasoning, clinical testing, and user interface. It should be pointed out, however, that from the beginning the system was designed to be integrated in an actual clinical setting. For this reason, emphasis is placed on automating certain aspect aspects of the decision-making process, while user interaction has been emphasized in those aspects of the

process where the clinician should, and probably prefers to, become involved in the process. This has recently resulted in an interface that both facilitates and invites usage. Another recent and important modification has been in terms of the implementation environment, which provides for greater access and portability. This paper describes these aspects of PERFEX, as well the methods employed in its development, implementation, and testing. In Section 2, a discussion is provided of the tomographic acquisition procedure, including an explanation of the importance of this procedure in terms of health care, and the type of information that the procedure generates. Section 3 is devoted to the knowledge representation, reasoning, and inferencing mechanisms, including the feature extraction process from which symbolic representations are obtained, the handling of uncertainty, and the integration of image, non-image, and temporal information. Section 4 is concerned with the implementation and testing of the system. Finally, Section 5 makes concluding remarks and summarizes some of the current research activities.

SECTION 2. MYOCARDIAL 3D IMAGING: METHODS AND RATIONALE

In order to better understand the methods employed in interpreting and visualizing 3D imagery, it would be worthwhile to first define the objectives of the underlying research and formulate the problems to be solved. In the broadest of terms, the overall objective is to provide a clinically useful, computer-based system to assist in diagnosing heart disease. Patients who are at risk of having heart disease generally undergo a series of clinical tests, including a nuclear "scan" that is performed to determine how well blood is perfusing throughout the myocardium (heart muscle). This test gives a measure of the extent and severity of disease, particularly in the left ventricle. The specific test is a cardiovascular nuclear medicine tomographic perfusion study, which is routinely used in medical centers. More precisely, the intent of this nuclear study is to detect and localize infarction (diseased or dead muscle tissue) as well as ischemia (viable myocardium but at risk of infarction). The procedure involves injecting the patient with a radiopharmaceutical tagged to go to the heart, such that the emitted radioactivity can be imaged, and, in turn, the image gives the myocardial perfusion distribution information.

Tomographic perfusion imaging using thallium-201 (henceforth referred to as Tl) is a clinically useful, noninvasive technique used to perform these studies [4]. Patients undergoing Tl perfusion imaging are exercised in a fasting state according to a multistage treadmill protocol. At peak exercise, 3.5 millicurie (mCi) of Tl are injected, with exercise continuing for an additional 60 seconds. After approximately 5 minutes, and again 3 to 5 hours later, images of the patient are acquired using a rotating large field-of-view camera. These images taken at different times constitute the stress and delay images, corresponding respectively to perfusion distributions while the patient is undergoing exercise and while he/she is at-rest. The stress image is useful in

determining perfusion defects in general (and infarctions in particular), while the delayed images are useful in determining possible redistribution of perfusion (i.e., ischemia, where a defect found during stress "disappears" or reverses when the patient is at rest).

For a tomographic acquisition, thirty-two images are acquired over a 189-degree arc, for a preset time of 40 seconds per image acquisition. These images are corrected for field nonuniformity and misalignment of the mechanical center of rotation. Following these corrections, the images are filtered and backprojected to reconstruct 6mm-thick transaxial tomograms that encompass the entire heart. Coordinate transformations are subsequently applied to the transaxial tomograms to generate oblique, long- and short-axis tomograms, parallel to the long and short axes of the left ventricular myocardium. Figure 1 shows the long- and short-axis slices displayed using a color table that maps those areas containing relatively high concentrations of radioactive tracer into brighter colors (e.g., yellow and gold), while those areas that contain relatively low tracer concentrations are mapped into darker colors (e.g., magenta and blue). For each tomographic cut or slice, a maximal-count circumferential profile is then generated. Each point of these profiles represents the maximum (radioactivity) counts per pixel along a radius that extends from the center of the left ventricle to the limit of the radius of search. The profile is constructed from the values along 40 equally spaced radii (i.e., every 9 degrees). This procedure is performed for each stress (patient undergoing exercise) and delayed (patient at rest) tomographic study. After all the circumferential profiles are extracted, they are then interpolated to represent 15 slices (myocardial "height") by 40 angular locations (myocardial circumference"). These 15-by-40 matrices are then plotted in a 2D polar representation, as shown in Figure 2, where the myocardial apex corresponds to the center region and the myocardial base corresponds to the periphery of the polar map. A similar plot (now shown) is generated for the delayed image.

INSERT FIGURES 1 AND 2 HERE

During a clinical study, images of each patient are compared with those of "normal" patients. For this comparison, normal limit profiles are generated using the mean and standard deviation established from pooled profiles of normal patients for each angular location and each depth in the tomographic study. For an individual patient, the stress and delayed profiles are compared with the corresponding, gender-matched normal profiles and, from these comparisons, 15-by-40 matrices are generated which quantitatively express the location and severity of abnormal regions. For the purposes of visual diagnostic interpretation, these 15-by-40 matrices are normally

displayed as color-coded polar plots in two forms: as standard deviation maps, representing the deviation from normal limits, and as a "black-out" map, in which those regions that fall beyond normal limits (usually 2.5 standard deviations from normal) are blackened out, as illustrated in Figure 2 (the standard deviation image is not shown in this figure). Thus, in the black-out displays, perfusion defects appear as contiguous pixels that have been set to zero (blackened out), whereas in the standard deviation image, each pixel that is below the mean normal value is converted to the corresponding number of standard deviations below that mean. Visual inspection is thus possible, since the color in either standard deviation or black-out image gives a measure of the degree of abnormality, with black representing severely abnormal. Similarly, the reversibility (i.e., delayed or at-rest) image can be visually analyzed to gauge the extent to which perfusion defects reversed after several hours: whereas in the black-out image, the black region correspond to hypoperfusion, a whitened (white-out) region in the reversibility map represents perfusion redistribution several hours after stress.

Through clinical experience, interpretation of this visual information becomes possible by recognizing the many different possible patterns in the images. Thus, clinical interpretation by humans involves understanding the meaning of these patterns in terms of several things: possible artifacts, the physiology of perfusion, and the relationship between these and possible ischemia, infarction, coronary vessel anatomy, other clinical data, and overall patient condition. The polar map shown in Figure 2 is illustrative of one of the more popular forms for performing diagnostic interpretation. Typically, expert clinicians view images such as those shown in Figures 1 and 2 to assess patient condition. It is evident from these images and the foregoing discussions that much work has been done on the quantification of myocardial perfusion, and that these efforts have provided clinicians with a better understanding of the significance of perfusion defects [3], thereby facilitating the assessment of heart disease. However, the resulting body of knowledge does not easily lead to mechanisms for computerized decision-support. One obvious challenge consists of capturing and representing domain knowledge, which includes both visual and temporal knowledge for interpreting images taken at different times. Also, clinically useful decision-support would require the integration of additional information that can be either numeric or textual in nature. And, symbolic representations of data as well as knowledge should allow for inexactness. These considerations present interesting knowledge engineering challenges and are the subject of the remainder of the paper.

SECTION 3. KNOWLEDGE REPRESENTATION, REASONING, AND INFERENCE METHODS

A diagram representing the basic elements of PERFEX is shown in Figure 3 in terms of overall information flow. As shown in this figure, the input information consists of the image-derived information previously described (i.e., images corresponding to stress, standard deviation, and reversibility) as well as other patient-specific information. This section will describe how this information is used in PERFEX, and how the different components of PERFEX shown in Figure 4 relate to each other.

INSERT FIGURE 3 HERE

Symbolic Feature Extraction

The visual input to PERFEX consists of the standard deviation (SD), black-out (BO), and reversibility (RV) images described above. The information contained in these images represents the deviation from normal limits for a particular patient, where the black-out regions are those associated with severe hypoperfusion (usually beyond 2.5 standard deviations from normal limits). These images, then, actually represent the three-dimensional myocardial perfusion distribution in a quantitative form. Thus, the first task is that of mapping this visual information into symbolic form to infer diagnostic interpretations, essentially a vision task. For image processing purposes, grey-level intensity rather than color is used (since the color code is only intended to facilitate human visual inspection); the images are also cast in the form of 15-by-40 matrices (rather than in polar formats) without information loss. Initially, a search is conducted through the BO array for any pixel that has been set to zero (i.e., abnormal). This is followed by edge-hugging operations to isolate all other pixels set to zero that are also connected to the first 0-valued pixel found. This process identifies the first perfusion defect, and an entry is made in a list to keep record of this defect. In a systematic fashion, the entire BO array is searched to detect each and all perfusion defects, making the appropriate additions to the list. Upon termination of the search, information exists in a data file with the total number of defects and their location of each within the image. Once the defects have been detected and ordered in this fashion, a symbolic description is assigned to each defect. The description is defined spatially, in terms of regions of myocardial depth (basal, medial, proximal apical, and distal apical), and myocardial walls (septal, anterior, inferior, and lateral regions, as well as their pair-wise combinations). This scheme yields 32 possible regional descriptions (or descriptors for brevity) for localizing each perfusion defect, as shown in Figure 3. One defect can, of course, be composed of different descriptors. An example of a (small) defect with only one descriptor would be: "Antero_Septal_Medial", identifying a defect in the region at 2 o'clock. The localization of features in terms of descriptors is used for all three types of images

used by PERFEX (BO, SD, and RV images).

INSERT FIGURE 4 HERE

Representation of data uncertainty and defect severity

To represent the severity of the perfusion defects, the standard deviation map was used. Each descriptor location in a defect (e.g., Antero-Septal_Medial) is associated with a descriptor in the SD map in a corresponding location. This SD descriptor has a numeric standard deviation value which, in fact, is a measure of severity. For the purpose of inferencing, as will be discussed later, this severity value can be expressed in terms of a certainty factor, obtained from an empirically-derived function that maps SD values to certainty factor values, the latter falling between -1 and +1 according to the Certainty Factor (CF) Model [5]. This provides a mechanism for representing evidence confirming (if the CF value is positive) or disconfirming (if the CF value is negative) the relative degree of abnormality associated with a descriptor (or groups of descriptors making up a defect), as well uncertain evidence (if the CF value is close to zero) of abnormality. The descriptors that have all its pixel values in the normal range are thus set to $CF = -1.0$, identifying them as definitely normally perfused regions; those having any of its pixel values in the abnormal range (2.5 standard deviations or more below the normal mean) are assigned according to the formula $CF = 0.145SD - 0.163$, where SD is the average number of standard deviations below the mean normal response for pixels in the abnormal range.

A similar procedure is repeated for the descriptors associated with the delayed or RV images. This RV descriptor arrays are determined as the normalized difference between the delayed and stress TI-201 distributions. The reversibility array provides information as to whether a stress defect normalizes with time and is thus consistent with an ischemic event. These stress and reversibility descriptor files are used by the knowledge base for image interpretation.

This method of representing uncertainty, location and severity with respect to stress and reversibility images was selected after much experimentation with different models of symbolic representation and inexact reasoning. The method is straightforward yet extremely powerful. Indeed, it is a symbolic mapping that is semantically meaningful, since the descriptors represent mutually distinguishable spatial regions with medically useful labels; the descriptors also cover the myocardial region of interest in sufficiently small segments while admitting evidence for or against disease for images taken at different times.

Knowledge representation

The knowledge base of PERFEX is designed to relate the symbolic information just described, in conjunction with other, patient-specific data, to the presence (or absence) of coronary artery disease. The knowledge representation is achieved in terms of rules derived in a knowledge acquisition effort that spanned several years and involved careful analysis of several hundred actual clinical cases, conducted in collaboration with expert diagnosticians [2]. Rule-based knowledge representation and acquisition are well known paradigms [5,6], and although they remain active and important research areas, the present discussions will be primarily devoted to how these paradigms relate to PERFEX.

The knowledge base is organized in terms of knowledge islands, each of which contains knowledge about particular aspects of the overall interpretive process. Figure 5 is an illustration of the organization of the knowledge base in terms of these knowledge islands. This figure also gives an indication of the hierarchy implicit in this organization structure and of the direction of flow of information. Since the knowledge structure is intimately connected with the inferencing mechanism, it may be worthwhile to describe both in a dynamic fashion, illustrated by means of a hypothetical consultation session.

INSERT FIGURE 5 HERE

The interpretive process begins by first considering patient-specific (non-image) information. The information that is first considered includes the patient's age, sex, symptoms, and electrocardiographic results. These are used to obtain an estimate of the pre-test likelihood of disease, i.e., the a priori probability of disease [7]. (Sex and age information may also be used by accessed by other rules during a consultation.) This pre-test likelihood is calculated deterministically, and is used in overall patient evaluation. Referring to Figure 5 once again, one of the initial interpretive tasks deals with the characterization of the image information in terms of size and location of perfusion defects. This is accomplished by using the symbolic descriptors and their corresponding CF values. An illustration of how rules use these descriptor values is:

RULE ILM LOCATION

```
IF          Defect_Descriptor   IS   ILM"

THEN          Defect_Location   IS   "Infero_lateral_medial"

AND_DO       Nothing_Now
```

FIGURE 6. Sample rule for assigning symbolic descriptions to myocardial regions of interest.

This example illustrates the fact that the rules are constructed in both declarative and procedural manner. In the example, the symbolic descriptor (ILM) led to the conclusion or declaration of the location of the perfusion defect (inferolateral region; see Figure 4), while at the same indicated what action or procedure should be undertaken at this point (nothing at the moment). Thus, the rules contain premise, conclusion, and action clauses.

At present, over 300 rules are contained in the knowledge base. The rules contain expert-based knowledge and reasoning strategies representing the diagnostic process. Following the structure illustrated in Figure 5, this knowledge is structured in terms of knowledge islands to perform the following interpretive tasks: (1) characterize the image information symbolically in terms of location, size, and shape of stress perfusion defects (as in the previous rule example); (2) relate this symbolic information to evidence that confirms (or disconfirms) disease in three major vessels: the left anterior descending (LAD), left circumflex (LCX), and right coronary artery (RCA); (3) recognize shapes that, although suggestive of multiple-vessel disease because the defects "spill" over several territories, they are actually associated with only one vessel); (4) recognize the presence of artifacts (i.e., determine whether a defect is actual or a result of such effects as improper tomographic slice selection, attenuation, etc.); (5) detect and localize regions of reversibility (the "opposite" of perfusion defects, denoted as "virtues" in Figure 5); (6) relate this temporally derived information regarding reversibility to specific arterial vessels (LAD, LCX, and RCA); (7) recognize "spillover" shapes associated with perfusion redistribution (i.e., reversibility); (8) recognize the presence of reversibility artifacts; (9) keep track of how multiple defects affect each of the three major coronary arteries of interest (LAD, LCX, and RCA); (10) suggest an overall diagnostic evaluation of the patient; and (11) provide a report that conforms to the standard format of the report that is routinely used clinically.

Another illustration of the type of knowledge representation used in these knowledge islands is given below:

RULE LAD-DEFECT-NOT-RCA

IF	Defect_Sector	IS	Septo_Inferior AND Septo_Anterior
AND	Diseased_Coronary	IS	Right_Coronary_Artery
AND	Defect_Sector	IS_NOT	Inferio_Lateral

```
THEN      Diseased_Coronary   IS      Right_Coronary_Artery (CF = -0.9)
DO      Decrease Right_Coronary_Artery_Evidence
```

FIGURE 7. A sample rule containing certainty factors and procedural clauses

Note that this is an example of a rule that considers sectors (rather than descriptors); these sectors represent larger myocardial regions and the symbolic representation in terms of sectors is built up from the descriptor, using rules similar to that shown in Figure 6s. In addition, the rule given in Figure 7 is an example of one that disconfirms evidence of disease in the right coronary artery (hence the negative certainty factor value at the lower right); this will be expanded on later.

Representation of Temporal Reasoning

As previously mentioned, stress perfusion defects may indicate the presence of disease. However, the delay image information can suggest that some or all of these perfusion defects are reversed at a later time, when the patient is at rest. Thus, the information associated with reversibility, which is captured in terms of 32 symbolic descriptors analogous to the stress descriptors, represents temporally defined information that modifies the conclusions that may have been inferred by PERFEX regarding possible disease, particularly regarding ischemia versus infarction. Thus, a temporal representation is necessary to account for these effects. It has been previously suggested that a full solution to the problem of time representation is an NP-hard issue, and requires certain compromises [8]. Although several logics of temporal inference exist [9, 10, 11], we have opted to follow a scheme similar to Rucker *et al.*'s of reducing complexity by using only specific time query representations, while adhering to the constraints associated with rule-based representation [8]. In our case, there is only one temporal phase of interest (associated with resting perfusion), thereby significantly reducing complexity. Thus, we employ a high-level representation of time in a two-phase procedure, corresponding to stress and delayed information. This scheme also resembles the reference intervals suggested by Allen [10] and the contexts of Kahn [9]. Therefore, the knowledge associated with reversibility is used as to modify inferences drawn from the stress imagery, where both phases are represented in an object-oriented framework. It should be added that most other applications involve multiple phases or segments of time, and that in those instances additional revision mechanisms may be necessary. In our case, however, the temporal representation is minimal (i.e., only two temporal segments need to be considered), and the revision mechanism occurs at a very high level to simply determine whether a particular defect is likely to be a "fixed" or "reversible" defect.

Inferencing and Uncertainty Reasoning

At this point, it is important to remark about the inferencing mechanism. During a typical consultation, PERFEX conducts a heuristic search that is primarily inductive in nature. As such, the inferencing is conducted in a forward chaining mode: from initial data to hypotheses (although it is possible to proceed in the reverse fashion). This is done primarily to incorporate the temporally based knowledge previously described. However, individual knowledge islands may be governed by locally deductive, goal-driven inferencing mechanisms.

As mentioned earlier, we have considered several models of reasoning with uncertainty in our research, including fuzzy set representation, probabilistic and qualitative reasoning, and belief networks [12]. However, it has been our experience that if the diagnostic accuracy obtained from different models are consistent with each other (in terms of sensitivity and specificity), then clinical utility will probably be decided by factors other than mathematical rigor (such as ease of use). With this in mind, we have selected those models that combine sufficient mathematical rigor with ease of implementation. At present, most of the uncertainty reasoning in PERFEX is based on the Certainty Factor Model, which is well known and has a number of well documented strengths and weaknesses [5]. In fact, it is precisely because of the wealth of knowledge that is available regarding the CF Model that it still remains in wide usage.

As stated earlier, the CF Model is used to express the relative certainty associated with the initial data, such as the severity of perfusion defects. In addition, this model is also used throughout the knowledge base. Thus, CF values are associated with the clauses in rules. The earlier sample rules, given in a simpler form for simplicity, more closely resemble the following form:

RULE ILM LOCATION

```
IF      Defect_Descriptor   IS      "ILM" (CF-ILM = 0.80)
THEN    Defect_Location     IS      "Infero_lateral_medial" (CF-LOC = 1.0)
AND_DO          Nothing_Now
```

FIGURE 8. A sample rule containing certainty factors.

In general, if the CF value of the premise of a rule is greater than 0.2, the rule can be fired. The CF values are propagated according to the algorithms developed in the CF Model [5]. In keeping with the model, the combinatorics allow for providing relative measures of belief as the inferencing proceeds toward the goals, which in our case are the hypotheses concerning the presence of disease in each of three coronary vessels, and an evaluation of overall patient condition, as suggested in Figure 5.

SECTION 4. SYSTEM IMPLEMENTATION, TESTING, AND RESULTS

A number of expert system shells have also been investigated throughout out research, including Intellicorp's KEE™, Texas Instrument's Personal Consultant™, IBM's ESDE™, and our own shell. At present, PERFEX is implemented in an object-oriented environment using Neuron Data's Nexpert Object™. This object-oriented framework provides a some advantages, including inheritance properties and C-code. This software, however, has been extensively modified to incorporate the CF Model (which is intimately linked to inferencing) and to allow for a dynamic user interface.

The interface consists of a screen that displays a patient report, as shown in Figures 9 and 10. Figure 9 shows the screen as it might appear during a consultation session, displaying the polar plots at the bottom portion of the screen and the patient report at the top portion. The interface allows the user to point (with a mouse) and click on any of the fields in the report for further exploration and justifications. A full patient report is also printed, shown in Figure 10. Although the report seems somewhat visually intricate, it contains all the pertinent clinical information that is normally associated with the report filled out manually by the clinician. Thus, the screen serves as a report-generation mechanism as well. In the actual consultation, the screen shows the original imagery as well as portions of this report which can be scrolled up or down. The interaction is by means of a mouse, such that the user can click a field within the report, and initiate the process of asking the system to justify the reasoning behind the results or conclusions shown in the selected field. At present, we are developing this interface to support X-windows and to also show additional explanatory text and annotations. The reason why this interaction paradigm is useful is that the traditional way of conducting a clinical study is by having the clinician sit at a workstation to visually interpret the imagery, and PERFEX is designed to be a (transparent) part of the workstation environment.

INSERT FIGURES 9 AND 10 HERE

A number of tests of PERFEX have been conducted or are underway. A preliminary pilot study consisting of a set of 20 patients has been conducted, comparing the interpretations made by PERFEX with those of a human expert in thallium image interpretation. The human expert interpreted that of the 20 patients, 3 were normal, 16 were consistent with coronary artery disease (CAD), and that one exhibited a count reduction in the anterior wall due to breast attenuation but

was also probably normal (i.e., an artifact). PERFEX agreed with all of these interpretations, including identification of the breast artifact, although it called the patient probably abnormal. Of the 16 patients with evidence of CAD, PERFEX agreed with the expert in 7 out of 8 patients with ischemia, and in all 8 patients with infarction (and no ischemia). In addition, all of the 28 vascular territories (LAD, LCX, or RCA) identified by the human expert were correctly localized by PERFEX, although the latter incorrectly identified an additional 10 abnormal vascular territories. Of the 28 abnormal vascular territories, PERFEX correctly classified 12 of 16 as reversible (ischemic) and 11 of 12 as fixed (infarcted or scarred). These results agree very favorably with experts' interpretations of the images. The other significant consideration is that the results require less than 30 seconds to be generated by PERFEX upon receiving the image information, and that the knowledge-based system resides directly embedded in the actual clinical environment. At present, an extensive study of 100 patients is underway, and multicenter trials are planned.

SECTION 5. CONCLUSIONS AND FUTURE DIRECTIONS

A knowledge-based approach to interpreting 3D tomographic myocardial perfusion distributions has been described. The approach is intended to suggest diagnostic recommendations regarding coronary artery disease. The methodology underlying this approach uses a rule-based paradigm to represent visual as well as temporal reasoning, and integrates patient-specific (non-image) information to provide more comprehensive diagnostic decision support. The approach provides a novel mechanism for mapping visual information into symbolic representations that allow for knowledge-based processing. Another innovative feature consists of inferring structural information (associated with coronary vessel anatomy) from physiological function (representing myocardial perfusion). The methodology can serve as a model for interpreting 3D image information in general, and other types of medical images in particular.

The approach has been implemented in an object-oriented framework which allows for user interaction to query the system about specific conclusions and recommendations. The implementation system, PERFEX (for perfusion expert) is presently undergoing a number of clinical tests to determine accuracy, reliability, robustness, and overall clinical utility. Preliminary results indicate that the diagnostic interpretations made by PERFEX are highly consistent with those made by human experts, and that several factors in its design enhance clinical utility, including a significant degree of automation coupled with the possibility of user interaction.

At present, plans are underway to extensively test and evaluate PERFEX. A 100-patient study has been completed, and the results are currently being analyzed. In addition, ways to implement the system in different nuclear imaging systems are being investigated, and multicenter testing is being planned. Another significant direction of research concerns the aggregation of knowledge regarding other types of relevant clinical information. Furthermore, efforts are underway to develop methods to visualize in three dimensions the tomographic information, thus facilitating the overall diagnostic process. It is envisioned that these current efforts will lead to a comprehensive way to interpret and visualize cardiovascular anatomy and function.

REFERENCES

1. R. W. Conners, C. A. Harlow, and S. J. Dwyer, "Radiographic image analysis: past and Present." Proceedings of the 6th. International Conference on Pattern Recognition, Munich, Germany, 1982.
2. N. F. Ezquerra and E. V. Garcia, "Artificial Intelligence in Nuclear Medicine Imaging," *Am. Journal of Card. Imaging*, **3**:130-141, 1989.
3. E. G. DePuey, E. V. Garcia, and N. F. Ezquerra, "Three-Dimensional Techniques and Artificial Intelligence in thallium cardiac imaging." *Am. J. Roent.* **152**:1161-1168, 1989.
4. E.E. DePasquale, A. C. Nody, E. G. DePuey, et al., "Quantitative rotational thallium-201 tomography for identifying and localizing coronary artery disease," *Circulation*, **77**:316-327, 1988.
5. B. G. Buchanan and E. H. Shortliffe, Rule-Based Expert Systems, Addison-Wesley, Reading, MA, 1984
6. Knowledge-Acquisition in Medicine, Lectures Notes, Tutorial at SCAMC (Symposium on Computer Applications in Medical Care), November 1990.
7. G. Diamond and J. Forrester, "Table 1: Likelihood after an Electrocardiographic stress test according to Age, Sex, Symptom, and Depression of S-T Segment," *New Engineering. J. of Med.*, **300**, pp.13501358, June 1979.
8. J D. Rucker, D. Maron, and E. Shortliffe, "Temporal Representation of Clinical Algorithms Using Expert-System and Database Tools," *Comp. and Biom. Res.* **23** 222-239, 1990.

9. M. Kahn, J. Ferguson, E. Shortliffe, and L. Fagan, "Representation and use of temporal information in ONCOCIN," Proc. 9th. Symp. Comp. App. in Med. Care, pp.17-176, Baltimore, MD, 1985.
10. J. Allen, "Towards a general theory of actions and time." Artif. Intl. **23** 123, 1984.
11. L. Fagan, "VM: Representing Time-Dependent Relations in a Medical Setting," Ph.D. thesis, Department of Computer Science, Stanford University, Stanford, CA, 1980.
12. For an up-to-date treatment of uncertainty reasoning, see, for instance, Lea Sombe, Reasoning Under Incomplete Information in Artificial Intelligence, J. Wiley & Sons, INC., New York 1990.