ROENTGEN: Case-based Reasoning and Radiation Therapy Planning*

Jeffrey Berger
Artificial Intelligence Laboratory
University of Chicago
Chicago, IL 60637

ABSTRACT

ROENTGEN *is a design assistant for radiation therapy planning which uses case-based reasoning, an artificial intelligence technique. It learns both from specific problem-solving experiences and from direct instruction from the user. The first sort of learning is the normal case-based method of storing problem solutions so that they can be reused. The second sort is necessary because* ROENTGEN *does not, initially, have an internal model of the physics of its problem domain. This dependence on explicit user instruction brings to the forefront representational questions regarding indexing, failure definition, failure explanation and repair. This paper presents the techniques used by* ROENTGEN *in its knowledge acquisition and design activities.*

INTRODUCTION

ROENTGEN is an artificial intelligence research effort with the practical goal of developing a computerized design assistant for radiation therapy planning (RTP). ROENTGEN, the program, has the following characteristics:

- It suggests suitable therapy plans for a given patient, points out possible problems with a plan under consideration, explains the cause of failures when they are observed, suggests repairs to correct a failed plan, and so on. It also has the ability to design therapy plans autonomously.
- ROENTGEN learns therapy planning by remembering successful plans. It functions as an assistant and designs its own plans by being reminded of these successful plans.
- ROENTGEN also learns from explicit instruction. The human expert can tell the system what features of a patient affect the design of a plan. ROENTGEN will take these features into account when looking for plans for future patients.

The kind of behavior we want for the system and the nature of RTP have led us to build ROENTGEN as a case-based reasoner[10]. Case-based reasoning (CBR) is an emerging research paradigm in AI. Researchers have investigated the power of this paradigm in various domains including planning[4], law[1], design[7], medical diagnosis[2, 8] and many other fields[5].

CASE-BASED REASONING

The primary source of knowledge for the CBR system is its memory of past problem-solving experiences.

When solving a new problem, the CBR system retrieves a similar case from its case memory. Next, the system adapts the solution from the retrieved case to account for any important differences between the new problem and the retrieved one. Then the system applies this suggested solution to the current problem. If this solution works, the system is done. If this solution fails, the system repairs it so that it correctly solves the new problem. Finally, the system stores the new solution back in its case memory in such a way that the failure will be avoided in similar circumstances. In this way, the system learns from its own problem solving experience.

The motivation behind CBR is that it is often easier to repair a nearly-correct solution than to build a fresh one from first principles. The paradigm takes advantage of the regularity which exists in many natural problem areas: small changes in the features of a problem usually result in small changes to the solution. ROENTGEN uses a case-based approach in order to take advantage of this regularity in the domain of radiation therapy planning.

CASE-BASED RADIATION THERAPY PLANNING

ROENTGEN's basic process diagram is shown in Figure 1. ROENTGEN's design cycle is: 1) from a patient description and prescribed dose to the target, retrieve a similar past case; 2) adapt the plan from the case to the current patient; 3) detect faults which occur when the proposed plan is simulated; 4) check for and correct errors in system knowledge; 5) determine if the plan is successful, repairing it if it isn't; 6) store the successful plan for future use.

The *Retriever*

ROENTGEN's first task, when it starts designing a plan for a new patient, is to remember a similar patient it has seen before. This is the *Retriever*'s job.

In order to find similar cases, the *Retriever* relies on the fact that therapy plans can be classified into types

*This work was supported in part by: DARPA contract number F49620-88-C-0058; DARPA contract number N00014-91-J-4092 monitored by the Office of Naval Research; Office of Naval Research grant number N00014-91-J-1185; and a grant from the Whitaker Foundation.

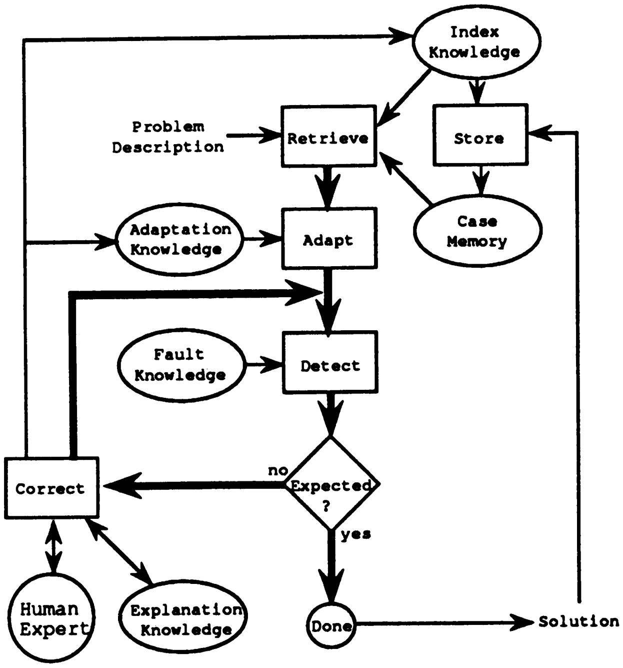

Figure 1: System diagram

according to the number and arrangement of their constituent beams. Moreover, each type has its own set of patient/tumor geometry preconditions and causally significant feature dimensions.

For example, in an opposed beam plan, two beams of photons are directed into the patient from opposite positions on the patient's periphery. These plans require as a precondition that there be a corridor through the patient's body which contains the target but does not contain any vital, radiation-sensitive tissue. Unless this precondition is met in a patient, it is unlikely that an opposed beam plan can be successfully used.

A causally significant feature dimension for opposed beam plans is "target depth ratio". Variation along this dimension requires changes in the relative weights of the two beams in the plan if it is to continue to perform satisfactorally. The feature is defined to be the ratio of the distance from the skin surface to the target center and the distance from the skin surface through the target to the skin on the opposite side of the patient.

When the *Retriever* is presented with a new patient, it tests for the applicability of each type of plan known to the system. If the patient/tumor geometry satisfies the preconditions for a type of plan, the causally significant feature dimensions for that type are used to find the best match[1] among cases with plans of that type in Case Memory. From all these best matches, the *Retriever* then selects the one with the highest match score. If there is a tie, the case with the simplest plan—fewest beams, fewest different beam energies, etc.—is selected. ROENTGEN will use the plan from the selected case in the design stages that follow.

[1] See section **The Storer**, for a more detailed description of how the match is computed.

The *Adapter*

Since there are sure to be differences between the patient for whom the plan is being designed and the one retrieved from memory, the retrieved plan will have to be adapted to the exact details of the new patient. The *Adapter* is the module charged with this task.

The *Adapter* adjusts the retrieved plan to account for differences between the retrieved and current patients using its adaptation knowledge. It is concerned with the same feature dimensions which the *Retriever* used to find a good match in memory. Recall that each such feature dimension has an effect on some plan parameter for plans of this type. This effect is represented by a table of feature value/parameter value pairs from the prior patients.

For example, suppose the system has learned that target depth ratio affects the relative weights of the two beams in an opposed beam plan. Suppose further that the system has seen two patients in which this affect was exhibited: one with a target depth ratio of .50 and relative beam weight ratio of 1.0 (the beams are left on for the same amount of time), and another with a target depth ratio of .32 and relative beam weight ratio of 1.5 (the second beam is on two thirds as long as the first beam). These value pairs will be in a table indexed under the opposed beam plan type in ROENTGEN's Adaptation Knowledge base. When working with an opposed beam plan, the *Adapter* will check the actual value of the target depth ratio in the current patient and use the table to interpolate the corresponding value of the relative beam weights. It sets each parameter for which it has a table to its interpolated value. So, ROENTGEN proposes the retrieved plan with some of its parameters adjusted by interpolation as a plan for the new patient.

The *Detector*

We now have a proposed plan for the new patient. The *Detector*'s job is to determine what the results of applying the plan will be.

A human designer uses a computerized dose calculator to determine the results of applying a treatment plan. The dose calculator takes the specification of a plan and the cross section of the patient and calculates the dose distribution the plan will produce. The information in the array is displayed as a contour map superimposed on the image of the patient cross section. The designer looks at this map to determine whether the target is receiving the prescribed dose, whether there are any vital tissues which are receiving too much radiation, and whether any other seriously undesirable conditions exist.

The *Detector* uses the results of the dose calculator as well. But it needs some way to interpret the significance of the information contained in the dose array and the tissue contours of the patient cross section. This knowledge resides in the Fault Knowledge base. The expert user builds up this knowledge base by providing definitions of fault conditions as they have occurred in patients the system has seen. Some faults warrant rejecting the plan that produces them. Others are not so serious; plans producing them may still be deemed acceptable.

The definitions of fault conditions are written using a language of primitives which allows the user to specify regions of the cross section corresponding to a particular tissue, refer to the dose deposited there, express whether or not that dose exceeds the limit for the involved tissue, and so on. All the conditions needed to make a judgement about the acceptability of a treatment plan can be expressed in this language.

The *Detector* looks for all the fault conditions that have been defined for the type of plan being used and the region of the body in which the treatment is being given. It produces a list of all the faults it discovers.

Expectation Failure?

ROENTGEN will be unable to contribute positively to the design process if the knowledge it's using in the given situation is in error. So, the system needs to determine whether the knowledge used on this patient is correct. Now that the results of the proposed plan are known, it can make that determination.

When ROENTGEN applies a plan adapted from a case in Case Memory, it has expectations about what the results will be. These expectations are simply the list of faults the *Detector* uncovered when the retrieved plan was applied to the retrieved patient. Because the new patient is similar to the retrieved patient and the current plan is similar to the retrieved plan, ROENTGEN expects that the results will be similar too.

If the results of applying the proposed plan are as expected, ROENTGEN assumes its knowledge is adequate to handle the current patient. If, on the other hand, the results are not as expected, ROENTGEN assumes its knowledge is deficient in some way.

The *Corrector*

If ROENTGEN experiences an unexpected fault, its knowledge is deficient in some way. The *Corrector*'s job is to correct ROENTGEN's knowledge so that the unexpected fault is explained. While this is accomplished, the plan will be repaired so that it no longer produces the unexpected fault.

A fault can be unexpected in several ways. One way is when the patient on which the current plan was originally used and the new patient differ along some causally significant feature dimension, but ROENTGEN is currently ignorant of this feature and its effects. The fault is being produced because there is a relationship between the feature dimension and some parameter in the plan which must be maintained if the fault is to be avoided, and the particular parameter value is not correct in the current plan. Generally, this sort of problem will occur when ROENTGEN has relatively little experience. All previous patients which were treated with the current type of plan must have had essentially the same value on this feature dimension, and hence were able to use the unchanged parameter value from the original plan of this type.

In this situation, the *Corrector* must correct ROENTGEN's Index Knowledge and its Adaptation Knowledge. The *Corrector* asks the human expert to provide a definition of the feature dimension of which ROENTGEN was unaware. She does this using the language of primitives which is part of the system. The *Corrector* adds this new feature dimension to the dimensions used by the *Retriever* when looking for plans of the current type. In this way, ROENTGEN learns new significant features for memory indexing.

Next the expert is asked to repair the plan, setting the maladjusted plan parameter to its correct value. With the corrected value, the *Corrector* can make a new entry in the Adaptation Knowledge base. The *Corrector* computes the feature value for the new patient and the original one. The *Corrector* gets the affected parameter value from the original plan and the corrected value for the same parameter in the proposed plan. It makes a feature/parameter table for the *Adapter* and places the feature/parameter value pairs from the two patients in the table. The table is placed in the Adaptation Knowledge base where it will be used by the *Adapter* for future patients. Then the *Corrector* adds an entry to its Explanation Memory which links the fault it didn't expect, the type of plan being used, and the feature/parameter table it just built. This packet of information is used to explain the causes of faults to novice therapy designers.

Finally, note that at the completion of processing by the *Corrector*, the user will have repaired the plan so that it no longer produces the unexpected fault that caused this module to become active. The repaired plan can be passed back to the *Detector* to make sure other faults were not introduced by the repair process.

The *Storer*

We are now at a point where a successful treatment plan has been developed for the patient. If the system can remember this plan and the circumstances under which it is appropriate, in the future it may avoid the work that went into making the plan. This is the job of the *Storer*.

The *Storer* adds new cases to memory, but only when 1) the plan has been approved by the human designer, and 2) the case required an extension of ROENTGEN's RTP knowledge. If ROENTGEN has learned about a new causally significant feature dimension, then the *Storer* builds a case from the new patient, plan, and plan results. By storing this case so that it can be found when needed by the *Retriever*, ROENTGEN and the designer will avoid having to repeat the effort involved in producing this plan.

The *Storer* adds the new case to the segment of Case Memory containing the cases with plans of the new case's type. Next, the *Storer* must update the Index Knowledge used by the *Retriever*. First, if ROENTGEN learned a new, causally significant feature while making the plan for this patient, the new feature must be added to the system's Index Kowledge. Second, the *Storer* must update the knowledge even for already-known features.

As cases are added to memory, how closely two patients match on a given feature dimension changes. Patients which were close enough on a feature dimension to be good matches when there were fewer cases should not match as well when cases which lie between them on that dimension are added to memory. Hence, the *Storer* alters the system's Index Knowledge to reflect these changes.

Once this is done for all the causally significant feature dimensions of the plan type, the *Storer*'s job will be complete. The *Retriever* will use this new scheme to retrieve cases in the future.

DISCUSSION

ROENTGEN as Advisor

The preceding description of ROENTGEN's modules is focussed on the early stage of the system's life when it is primarily concerned with learning new types of plans, and new features which are important when retrieving and adapting plans. Once it has accumulated enough of this sort of information, it will be able to carry out its role of advisor or critic. The *Retriever* and *Adapter* can combine to suggest candidate plans to a novice designer. The *Detector* can be run on the results of a plan to point out faults which the human designer may want to correct. The Explanation Memory can be used to determine the cause of faults produced by a plan the novice has developed.

Related Work

One of the main characteristics which distinguishes the CBR approach from other AI paradigms is the role and nature of domain knowledge. CBR efforts have been aimed at producing reasoning systems which have much less *explicit* knowledge about the domains in which they work than do traditional AI systems. To construct complete, explicit models of even relatively simple real-world domains is an extremely time-consuming, if not ultimately impossible, task. The case-based paradigm seeks to avoid this hard work while still producing powerful systems by focusing on *remembering* rather than *creating* solutions. ROENTGEN is unique in case-based research for the degree to which it shuns even the small amounts of explicit domain knowledge built into such systems. It relies on the human expert to tell it what features are important for memory indexing, what faults in plan results look like, how to repair problematic solutions and so on. As such, ROENTGEN represents a further advance within case-based research.

ROENTGEN differs from previous AI work in RTP[9, 6] in several ways. Its knowledge is memory-based rather than rule-based. Hence, it will learn new plans through its experience and become more capable at its task over time. ROENTGEN is an interactive system as well as a stand-alone one. It can support the human designer with suggestions throughout the design task and ask the designer to supply knowledge it needs to solve a problem. This knowledge then becomes a permanent part of the system's knowledge. Finally, ROENTGEN doesn't attempt to produce an "optimum" plan. The question of whether or not a plan is acceptable is left to the human designer. This is a strength, not a weakness. Researchers in RTP have not yet reached a consensus on an objective definition of plan optimality[3].

Future Work and Evaluation

Future work. Geometry is all-important for ROENTGEN. In our approach, we've assumed that if two patients have similar geometry and similar dose-to-target prescriptions, successful plans for them will also be similar. But, other factors may also influence the physician's therapy plan choice. To a certain extent, these other factors might be translatable into changing the dose tolerances for sensitive tissues in the patient. The faults the Detector looks for are functions of these tolerances. Hence, if the patient has poor pulmonary function, explicitly lowering the lung tolerance for that patient would cause ROENTGEN to be relatively more protective of the lungs. However, this would only effect the current design effort. Some way of representing these other factors would still have to be worked into the memory indexing machinery.

While geometry is all-important for ROENTGEN, it currently looks at the geometry of a single two-dimensional section of the patient. Most external beam radiotherapy currently being provided is two-dimensional, i.e., all the beams in a plan lie in a single plane. But, three-dimensional planning is an active area of research in RTP. Planning in three dimensions is much more complicated and less well understood than in two. Because of the much larger number of plans possible in three dimensions, a computer-based assistant should be that much more of an aid. An area of future work is to extend ROENTGEN's abilities to three-dimensional RTP.

Evaluation. As for evaluation, there are three questions we would like to have answers to:

- At what level does ROENTGEN perform when functioning autonomously with a mature case-base? Are its plans distinguishable from those of an expert? A novice?
- How does ROENTGEN affect the performance of a novice or expert designer when it functions as an assistant to the human?
- What is ROENTGEN's effect on the work of a treatment center? Are different sorts of plans produced at centers which use ROENTGEN as opposed to ones that don't? Is the amount of time to produce plans less? What is the effect on training time?

We will need more experience with ROENTGEN in a clinical setting to devise a reliable and convincing test of its effects on practice.

CONCLUSION

ROENTGEN is a case-based system which works on RTP problems. It learns by remembering therapy plans as they are designed by human experts and by explicit instruction from the human user. ROENTGEN has no built-in knowledge of the RTP domain and relies on communication with the human user by way of vocabularies of primitives which can be used to define success preconditions for types of plans, causally important features, and faults in plan application. After building up its case memory and other domain knowledge, it can function as an autonomous therapy planner, or an assitant which provides suggestions or points out problems to a human planner. ROENTGEN's case-base and domain knowledge can also serve to pool the expertise at a treatment center so that the experience of the most skilled therapy planners is available to the entire planning staff.

Acknowledgments

Kristian J. Hammond, Artificial Intelligence Laboratory, University of Chicago, provided essential suggestions and guidance for my work on ROENTGEN. George T. Y. Chen and other members of the Department of Radiation and Cellular Oncology, Billings Hospital, University of Chicago, have generously tried to educate me about radiation therapy planning. In addition, Timothy Converse, Patricia Goldweic, Mitchell Marks, Thomas McDougal, and other members of the Artificial Intelligence Lab, have been patient listeners and persistent questioners.

References

[1] K. Ashley and E. Rissland. Compare and contrast, a test of expertise. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 273–284, Palo Alto, 1987. AAAI, Morgan Kaufmann, Inc.

[2] Ray Bareiss. *Exemplar-Based Knowledge Acquisition*, volume 2 of *Perspectives in Artificial Intelligence*. Academic Press, San Diego, CA, 1989.

[3] Micheal Goitein and A. Niemierko. Biologically based models for scoring treatment plans. In S. Zink, editor, *Future directions of computer-aided radiotherapy*, San Antonio, TX, 1989. National Cancer Institute.

[4] K. Hammond. Chef: A model of case-based planning. In *The Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, PA, August 1986. AAAI.

[5] K. J. Hammond, editor. *Proceedings of a Workshop on Case-Based Reasoning*, Palo Alto, 1989. Defense Advanced Research Projects Agency, Morgan Kaufmann.

[6] Ira J. Kalet and Witold Paluszyński. A production expert system for radiation therapy planning. In Allan H. Levy and Ben T. Williams, editors, *AAMSI Congress*, pages 315–319, Washington, D.C., 1985. American Association for Medical Systems and Informatics.

[7] J.L. Kolodner. *Retrieval and Organizational Strategies in Conceptual Memory*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1984.

[8] Phyllis Koton. *Using Experience in Learning and Problem Solving*. PhD thesis, MIT, 1988.

[9] Witold Paluszyński. *Designing Radiation Therapy for Cancer: An Approach to Knowledge-based Optimization*. PhD thesis, University of Washington, 1990.

[10] R.C. Schank. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge University Press, 1982.