Integrating Case-Based and Causal Reasoning

Phyllis Koton[1]
Laboratory for Computer Science
Massachusetts Institute of Technology

Introduction

Much research in artificial intelligence has been directed toward the development of high-performance domain-specific problem solving systems, called *expert systems* or *knowledge-based systems*. Many current expert systems rely on *associational* knowledge (heuristics, empirical associations, "rules of thumb") for their reasoning power. An alternative approach, *causal reasoning*, uses a model of the objects in the domain and the operations that can be performed on those objects. Causal models can provide richly-detailed knowledge bases for reasoning in many domains, but models are inefficient compared to the associational knowledge typically used in expert systems.

Human problem solvers are able to use both associational and causal reasoning. We recognize and quickly solve common problems, but can use more detailed causal knowledge when faced with novel or difficult problems. An artificial reasoning system that combined both types of knowledge, using associational knowledge for speed, and reserving the ability to reason from a model when necessary, similarly, would be highly desirable. Case-based reasoning techniques [Kolodner, 1985] can be used to improve the performance of causal model-based systems, because the act of retrieving a similar case and using its solution is clearly associational: features of a problem are associated with a solution to that problem. However, case-based reasoning techniques have not been widely used to enhance the performance of expert systems. Also, case-based reasoning systems suffer from lack of a model. There is no justification (other than coincidence) for believing that transferring a solution from a previous case to a new case will produce a valid solution. The combination of case-based reasoning techniques with causal reasoning could result in substantially improved expert systems.

Overview of CASEY

I have developed a program, CASEY, which integrates case-based and causal reasoning. The causal reasoning component employs a model of the cardiovascular system developed for the Heart Failure program [Long, et al, 1986], an expert system for managing patients with heart disease. The case-based reasoning component uses a self-organizing memory system [Kolodner, 1983] to store descriptions of all patients the program has seen, and generalizations derived from similarities between the patients. The patient description is comprised of *features*. These include both input data, such as signs and symptoms, test results, history and current therapy information, and solution data, such as the causal explanation for the patient, the diagnosis, therapy recommendation and outcome information.

CASEY's output is a causal explanation that describes a relationship between physiological states in the model and observable features of the patient. This is produced using a five-step

[1]The work reported here has been supported (in part) by National Institutes of Health grants R01 LM 04493 from the National Library of Medicine and R01 HL 33041 from the National Heart, Lung, and Blood Institute. Robert Jayes, MD kindly provided the example cases. William Long's Heart Failure program provided an excellent resource for this work. Peter Szolovits, Ramesh Patil, and William Long gave much helpful advice in supervising this research.

process. First, CASEY finds a case similar to the new patient in its case memory. Next, it evaluates the significance of any differences between the new case and the retrieved case. During this phase the match can be invalidated if there are significant differences. If none of the differences invalidate the match, CASEY adapts the solution from the retrieved case to fit the new case. If a match is ruled out, or if no similar previous case is found, CASEY uses the Heart Failure program to produce a solution for the case *de novo*. The new case and its solution are stored in CASEY's memory for use in future problem solving.[2] Finally, the features which were causally important in the solution of this problem are noted in the memory.

Retrieving, adapting, and storing cases are standard procedures of a case-based reasoner. Because the match between a new problem and a previously solved problem usually is only partial, there may be differences between the two cases that preclude using even a modified version of a retrieved solution for a new problem. The justification step ensures that a retrieved solution can be supported by the data in the new problem. Evaluating the features of the new problem to determine which were important to the solution helps the program make better matches in the future, because it allows the program to distinguish between random features and important ones.

CASEY differs from previous case-based reasoning systems because it uses information from its causal model to effect these steps. During retrieval, causal knowledge is used to select important features of the new case for matching. During justification, causal reasoning is used to judge the significance of differences between the new and previous cases. In repairing a retrieved solution, the causal model determines what changes should be made in the retrieved solution so that it fits the present case. Feature evaluation uses the causal explanation of the new case to determine its important features.

Matching and Retrieval

Most case-based reasoners use a similarity metric to gauge the similarity of two problems, and to choose the best match for the current problem. Similarity metrics use a combination of the number of features in common and the relative importance of those features. In many case-based reasoning systems, the relative importance of features for matching is predetermined by the system designer (for example, [Simpson, 1985], [Hammond, 1986], [Bain, 1986]). Another approach is to determine the important features based on a program's experience of what was important in solving similar problems.

CASEY matches a new case against cases in its memory using *every* feature in the patient description. However, not all the features are equally important in matching a case to a previous case. Furthermore, the important features for matching may vary from case to case. Therefore, CASEY determines the important features for matching dynamically for each new case presented to the system, and gives these features greater weight for matching. Important features are defined as those that played a role in the causal explanation of previous similar cases.

Figure 1 shows a sample patient presented to CASEY, Oprah. The retrieved case for Oprah, a patient named Mary, is shown in Figure 2.[3] The features marked with an asterisk in Mary's description are those that were used in the solution to her case. Oprah has some, but not all, of the features that were important in the case of Mary. The two cases also share some features that were

[2] The user has the option of rejecting CASEY's solution, in which case Heart Failure program is used to produce a causal explanation, which will be stored in memory.

[3] The patient descriptions in these illustrations have been simplified to conserve space by excluding features with normal values. A patient description typically consists of about 40 features.

```
(DEFPATIENT "Oprah"
HISTORY                    PHYSICAL-EXAM
(AGE . 77)                 (APPEARANCE NO-APPARENT-DISTRESS)
(SEX FEMALE)               (PULSE SLOW-RISE)
(DYSPNEA ON-EXERTION)      (AUSCULTATION MURMUR S2)
(CHEST-PAIN ANGINAL)       (MURMUR SYSTOLIC-EJECTION-MURMUR)
(ANGINAL UNSTABLE)         (S2 SOFT-A2)
VITAL-SIGNS                (APEX-IMPULSE SUSTAINED)
(BLOOD-PRESSURE 147 89)    LABORATORY-FINDINGS
(HEART-RATE . 84)          (EKG LVH NORMAL-SINUS)
(RESP . 14)                (CXR CARDIOMEGALY)
(TEMP . 98.6)              (CARDIOMEGALY LV))
```

Figure 1: Description of patient Oprah.

```
(DEFPATIENT "Mary''
HISTORY                               PHYSICAL-EXAM
(AGE . 67)                            (APPEARANCE DIAPHORETIC* ANXIOUS*)
(SEX FEMALE)                          (PULSE NORMAL)
(DYSPNEA ON-EXERTION*)                (AUSCULTATION MURMUR S2)
(CHEST-PAIN ANGINAL)                  (MURMUR SYSTOLIC-EJECTION-MURMUR*)
(ANGINAL UNSTABLE* EXPERIENCING*)     (S2 SINGLE*)
VITAL-SIGNS                           (APEX-IMPULSE SUSTAINED*)
(BLOOD-PRESSURE 148 90)               LABORATORY-FINDINGS
(HEART-RATE . 99)                     (EKG NORMAL-SINUS LV-STRAIN*)
(RESP . 14)                           (CXR CARDIOMEGALY)
(TEMP . 98.7)                         (CARDIOMEGALY GENERALIZED*))
```

Figure 2: Description of patient Mary.

not used in the solution of Mary's case. CASEY matched Oprah to Mary using the *similarities* between the two cases. It now calculates the *differences* between the two cases (shown in Figure 3) and passes this information to the justifier, along with the solution retrieved from Mary's case.

Justification and Adaptation

Two cases might have many similar features yet have one critical difference that invalidates the match. The critical question is whether different values of features in the problem description still support the same solution. In CASEY's domain, this manifests itself as whether different patient symptoms still support the same causal explanation. CASEY therefore uses a set of *evidence principles* to evaluate differences between the new case and a retrieved case by examining the relationships between evidence and physiological states in the Heart Failure model. These principles rely on such concepts as alternate lines of evidence for states, additional supporting evidence for states, and inconsistent evidence. The module in CASEY that performs this evaluation is called the *justifier*.

A difference is *insignificant* if it does not affect the retrieved causal explanation. For example, the difference between Mary's and Oprah's temperature is insignificant because both temperatures are normal. A difference is said to be *repairable* if the features of the new case can be fit to the retrieved causal explanation. Consider, for example, the fragment of Mary's causal explanation shown in

Feature name	Value for Mary	Value for Oprah
age	67	77
anginal	unstable* experiencing*	unstable
blood-pressure	148/90	147/89
heart-rate	99	84
temp 98.7	98.6	
appearance	diaphoretic* anxious*	normal
pulse	normal	slow-rise
s2	single*	soft-a2
ekg	lv-strain*	lvh
cardiomegaly	generalized*	lv

Figure 3: Differences between patients Mary and Oprah.

Figure 4. The feature EKG: LV STRAIN is evidence supporting the state LV HYPERTROPHY. Oprah does not have lv strain on her EKG. CASEY looks in the causal model for other features that are evidence of LV HYPERTROPHY, and determines whether Oprah has any of those features. In fact, Oprah has two such findings: EKG: LVH and LV CARDIOMEGALY. Therefore, CASEY can justify keeping this state in Oprah's causal explanation. If all differences between the new case and the retrieved case are insignificant or repairable, then the transfer of solutions from the precedent to the current case proceeds.

In case-based reasoning systems without causal models, the problem solver finds the best match, transfers its solution to the new case, and hopes for the best. Sometimes a retrieved case leads the problem solver down the wrong path. Evaluating differences by use of a causal model improves the likelihood that the retrieved solution applies to the new case. When CASEY justifies the match between the old case and the new case, it demonstrates that although there are differences between the cases, the causal model still supports the retrieved solution.

Modifications to the solution are necessary for partial matches between cases. *Repair strategies* are invoked by the justifier when it discovers a repairable difference between the new case and the retrieved case. Repair strategies adapt a previous solution to a new case by adding or removing nodes and links on a copy of the retrieved causal explanation. In the example of Mary and Oprah, the justifier would invoke causal repair strategies to remove the evidence EKG: LV STRAIN from the state LV HYPERTROPHY, and add the evidence EKG: LVH and LV CARDIOMEGALY to the list of evidence supporting that state. Repaired solutions do not have to be tested (as is required in, e.g., [Hammond, 1986], [Simmons & Davis, 1987]) because the validation of the solution by the causal model has already taken place in the justification phase.

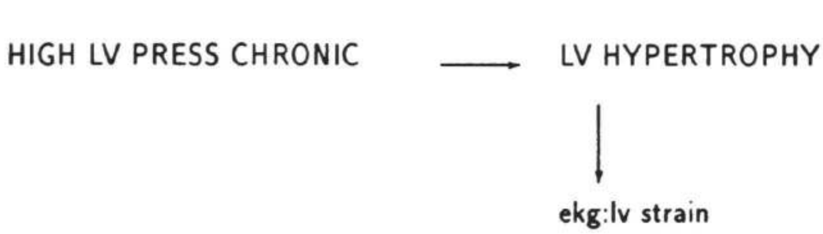

Figure 4: A fragment of Mary's causal explanation.

Storage

The use of causal reasoning influences the way a new case is stored in the case memory. A cases is indexed both by the input features that describe the case and the causal explanation that was derived for the case. CASEY also makes generalizations about the cases it has solved by finding similarities between the new case and cases already in its memory.

Generalizing the patient descriptions allows CASEY to make predictions about patients who share features [Kolodner, 1985] by recognizing co-occurrences. CASEY generalizes *all* the features in the patient description, not just the causally-related features. The Heart Failure model is incomplete, so it is possible that certain non-causal features are related to (and therefore can predict) some cause not represented in the model. The sex of the patient is an example: no state in the causal model uses the sex of the patient as evidence, yet there are causal relationships between gender and heart disease. The age of the patient is another example. The use of case-based reasoning therefore allows CASEY to *improve* on the performance of the Heart Failure system by learning new associations between features and solutions. At the same time, making generalizations about groups of similar patients reduces the effect of noise (random, unimportant features in the patient description) on the performance of the program. This is because spurious features are likely to occur randomly, whereas important features will tend to recur with some regularity in cases presented to the program.

Generalizing the explanations produces partial explanations that explain the features that the new case and retrieved case have in common. This allows CASEY to produce a partial solution for a similar problem in the future even in the absence of enough information to explain the whole problem.

Feature Evaluation

The use of a causal model is essential to CASEY's feature evaluation step. After producing the causal explanation for a new case, CASEY determines the important features of the *new* case (those features which were used as evidence in the causal explanation) and increments the weights of these features in the case memory. For example, after the case of Oprah is solved, the features EKG: LVH and LV CARDIOMEGALY are given extra weight in the memory because the causal model says that these features will be useful in identifying future patients with LV HYPERTROPHY, i.e. these features predicted a part of the solution. Determining the importance of features by experience is reasonable because the usefulness of a feature cannot always be determined in advance. This also allows the problem solver to adapt to changes over time in the types of problems is is presented. Giving extra weight to causally-related features is reasonable because because causality often indicates which features are important in the case for matching [Winston, 1981], [Schank, 1986].

CASEY also identifies the states in the causal explanation of the new case that are directly linked to findings, and stores the new case in memory using these states as indices. Future cases that contain evidence to support these states will retrieve Oprah's case as a match.

Related Work

CHEF [Hammond, 1986] combines case-based reasoning with a simple causal model. Therefore its causal reasoning can consist solely of chaining rules backward from an observed failure to a cause. This would not scale up to a reasonably sized domain. Recent work by Resnick and Davis [Resnick & Davis, 1988] combines a memory of past cases with explanation-based generalization of a causal model to produce a generalized description of a hardware fault, but their technique requires an exact match between the new problem and the old problem. IVY [Hunter, 1987], like CASEY, uses information of what was important to its reasoning task to select important features of the problem for storage. IVY, however, simplifies the selection problem by using heuristics to dispense with most of the features presented to the program.

Other systems have combined reasoning from a causal model and associational reasoning. ABEL [Patil, 1981] maintained a description of the patient at five levels of detail. It did not have a learning component. The Generate, Test and Debug method [Simmons & Davis, 1987] always uses its causal model to test proposed hypotheses, and always generates them using associational rules. CASEY *decides* when to use associational or causal knowledge. It determines whether the new case is sufficiently similar to ones it has already solved to use the associationally-derived solution. If not, it resorts to causal reasoning.

Discussion and Conclusions

CASEY *integrates* causal and case-based reasoning techniques in a program which is efficient, can learn from its experiences, and solves commonly-seen problems quickly, while maintaining the ability to reason using a detailed knowledge of the domain when necessary. The causal component is enhanced by the ability of the case-based component to learn new associations and compile detailed reasoning structures into simple associations between features and solutions. The case-based component is improved by the use of a causal model because the model can demonstrate that a retrieved solution will be helpful for a new case, and the model can be used to identify important features for matching.

Since determination of important features is based on information in the causal model, it is reasonable to ask why the Heart Failure model is not simply "compiled" to produce all this information in the form of associational rules relating important symptoms and physiological states. In fact, that is exactly what CASEY is doing, but it is compiling the knowledge incrementally, associating features of problems with solutions for the cases it has seen. Also, the Heart Failure program can generate solutions involving multiple diagnoses, its model provides the relative importance of features only for single diagnoses. To compile all of the Heart Failure program's knowledge taking into account multiple diagnoses would be computationally intractable. Because CASEY also makes generalizations about patients who have multiple diagnoses, it can create associational knowledge relating features to solutions involving multiple diseases.

CASEY's most serious limitations is that it assumes that there is very limited interaction between entities in the underlying causal model. For example, it has no concept of two states jointly causing a third. This weakness is due to the fact that no such interactions are represented in the Heart Failure model. This limits the generality of CASEY's reasoning techniques to other systems which also make this assumption. CASEY's evidence principles currently are being extended to handle more complex interactions.

Even if CASEY cannot solve a new problem completely because it lacks sufficient input information, it can often give a partial solution. Similarly, when CASEY is given a problem at the boundary of its expertise, it can produce a partial solution for features that are handled by its causal model, and leave the remaining features unexplained. Therefore, techniques such as are used in CASEY might be useful in integrating several causal models for different domains. Each model might use those features which it could explain, leaving the other features for other models. This is a topic for future work.

References

[Bain, 1986] William M. Bain. A case-based reasoning system for subjective assessment. In *Proceedings of the National Conference on Artificial Intelligence*, pages 523–527, American Association for Artificial Intelligence, 1986.

[Hammond, 1986] Kristian Hammond. *Case-based Planning: An Integrated Theory of Planning, Learning and Memory*. PhD thesis, Yale University, 1986.

[Hunter, 1987] Larry Hunter. Indexing with knowledge goals: how to put knowing what you know to use. In *Proceedings of the Forth International Workshop in Machine Learning*, 1987.

[Kolodner, 1985] Janet L. Kolodner. *Experiential Processes in Natural Problem Solving*. Technical Report GIT-ICS-85/22, School of Information and Computer Science, Georgia Institute of Technology, 1985.

[Kolodner, 1983] Janet L. Kolodner. Maintaining organization in a dymanic long-term memory. *Cognitive Science*, 7:243–280, 1983.

[Kolodner, 1988] Janet L. Kolodner. Retrieving events from a case memory: a parallel implementation. 1988. Submitted to AAAI88.

[Long, et al, 1986] W. J. Long, S. Naimi, M. G. Criscitiello, and R. Jayes. Using a physiological model for prediction of therapy effects in heart disease. In *Proceedings of the Computers in Cardiology Conference*, IEEE, October 1986.

[Patil, 1981] Ramesh S. Patil. *Causal representation of patient illness for electrolyte and acid-base diagnosis*. TR 267, Massachussetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, October 1981.

[Resnick & Davis, 1988] Paul Resnick and Randall Davis. Improving performance of model-based reasoning by learning from experience. March 1988. Submitted to AAAI88.

[Schank, 1986] Roger C. Schank, G. C. Collins, and Larry Hunter. Transcending inductive category formation in learning. *Behavioral and Brain Sciences*, 9:639–686, 1986.

[Simmons & Davis, 1987] Reid Simmons and Randall Davis. Generate, test, and debug: Combining associational rules and causal models. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 1071–1078, 1987.

[Simpson, 1985] Robert L. Simpson. *A Computer Model of Case-Based Reasoning in Problem Solving: An Investigation in the Domain of Dispute Mediation*. Technical Report GIT-ICS-85/18, Georgia Institute of Technology, 1985.

[Winston, 1981] Patrick H. Winston. *Learning New Principles from Precedents and Exercises: The Details*. AIM 632, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA, 02139, 1981.