Anthony Andrade

CS538

The workings of the CASNET model can be best described by dividing it into two sections. First I will describe the links that connect the nodes of the causal-association network. Then I will explain how they are used to interpret test results and suggest a therapy. There are three types of links in the first section, links between observations and p-states, links between p-states and disease-states, and links from disease state to disease state. These are the arc-values are provided by experts when building or refining the model. The second section will discuss the values used in interpreting observations/ test results and those used in selecting the next data to gather. These arc values are calculated at run-time and are changed when new data enters the system.

As mentioned earlier, the first section contains arcs from tests/observations to pathophysiological states, from these states to disease states, and from one disease state to another. The link from test to p-states is an associational link. Its weight ranges from –1 (deny) to 1 (confirm) and is used to represent the confidence that the p-state exists or doesn't exist in a patient given the observations. The next step up the model is an associational link from p-state to disease state. A classification table provides this association. The table is a conjunction of p-states that must exist or not exist in order to indicate a disease and an associated level of progression. Suggested treatments are represented in the same table by augmenting them onto the p-state-disease state combinations. The final foundational link type is that of the causal link from one disease state to another. This is where the temporal aspect of the CASNET model is represented. To follow the progression of a disease over time, the program simply moves from one disease state to the next. The arcs between disease states are weighted with a symbolic value based on the frequency at which one state leads to another.

The next section of the model describes what happens when a user interacts with the consultation program, rather than the model construction/refinement program covered by the first section. This section covers the strategies used for interpreting test results and selecting of new tests. The interpretation of test results is a calculation of how confident the system is that a particular result indicates a p-state. Once all of the changes in confidence are made through the effected p-states an evaluation of the indicated disease and its level of progress can be made. The weight that is used in this calculation is a confidence factor (CF).

Initially all of the confidence factors are set to zero, meaning that the system has neither confirmed nor denied any p-states yet. Different algorithms for determining the CFs have been used in CASNET. An example of one that was used would be to change the CF to the new CF, indicated by the test, only if it confirmed or denied the state with a greater magnitude than the old CF.

The other calculation made while the consultation program runs is the selection of new tests. This is doing the same kind of calculation as the interpretation, only now it is a prediction and so the calculation changes. One strategy for this calculation is as follows. Each test-node is given a weight. The weight given to the node represents a combination of the confidence, which the test will be able to indicate the p-state (the forward weight) as well as the p-state's ability to confirm a disease state (the inverse weight). By combining the weights going into each of the p-states with its ability to determine disease states the system can evaluate which p-state it is most useful to test for, and which tests will most confidently indicate that state. The cost of doing the calculation is also factored in after the weight calculation is made. The system can be changed to find the test with the highest cost-benefit ratio or look for the test with the highest benefit from within a range of costs.

At design-time, the expert sets up a system of links each with a weight to represent some degree of relationship between two nodes. At run-time the system uses the initial weights along with information from the user to calculate a disease-state and associated treatment. The system also makes a prediction about the significance of certain unknown information, and suggests tests to the user so it can make a more confidant diagnosis.