

Graphical Probabilistic Models in Diagnosis of Liver Disorders

Agnieszka Oniśko
Białystok University of Technology
Institute of Computer Science
Białystok, 15-351, Poland
aonisko@ii.pb.bialystok.pl

Marek J. Druzdzel
University of Pittsburgh
Decision Systems Laboratory
School of Information Sciences
and Intelligent Systems Program
Pittsburgh, PA 15260, U.S.A.
marek@sis.pitt.edu

Hanna Wasyluk
Medical Center of Postgraduate Education
Warsaw, 01-813, Marymoncka 99, Poland

12 May 1998

Our work concentrates on applying a graphical probabilistic model to the problem of diagnosis of liver disorders. Our work is continuation of the HEPAR project [1], conducted in the Institute of Biocybernetics and Biomedical Engineering of the Polish Academy of Sciences in collaboration with physicians at the Medical Center of Postgraduate Education. The HEPAR system includes a database of patient records from the Gastroenterological Clinic of the Institute of Food and Feeding in Warsaw, thoroughly maintained and extended with new cases. (The database available to us consisted of 570 patient records.) Our model is essentially a directed probabilistic graph modeling causal relations among a small set of essential domain variables with its numerical parameters extracted from the HEPAR database. The system is currently used in the clinic as a diagnostic and training aid.

There are several probabilistic graphical modeling tools (see for example [3]), of which the most popular are directed acyclic graph (DAG) models. One class of DAG models, widely used in the domain of Artificial Intelligence, is a Bayesian network [2] (also referred to as *belief network*, *probabilistic network*, or *causal network*). Each node of a Bayesian network graph represents a random variable and each arc represents a direct dependence between two variables. Formally, the structure of the directed graph is a representation of a factorization of the joint probability distribution. From the point of view of knowledge engineering, graphs that reflect the causal structure of the domain are especially convenient — they normally reflect expert's understanding of the domain, enhance interaction with a human expert at the model building stage and are readily extendible with new information. Quantification of a Bayesian network consists of prior probability distributions over those variables that have no predecessors in the network and conditional probability distributions over those variables that have predecessors. These probabilities can easily incorporate available statistics and, where no data are available, expert judgment. A probabilistic graph represents explicitly independences among model variables and allows for substantial savings in terms of both representation of the joint probability distribution over the model variables and in computational complexity of belief updating. Bayesian networks have been successfully applied to a variety of problems, including medical diagnosis, machine diagnosis, user modeling and user interfaces, natu-

ral language interpretation, planning, vision, robotics, data mining, and many others (for examples of successful real world applications of Bayesian networks, see March 1995 special issue of the journal *Communications of ACM*).

Our current network is based on data from the HEPAR database. Each of the records in the database is described by 119 features (binary, denoting presence or absence of a feature or continuous, expressing the value of a feature) and one of 16 classes of liver disorders. The HEPAR database assumes that a patient appearing in the clinic has at most one disorder. The features can be divided conceptually into three groups: symptoms and findings volunteered by the patient, objective evidence observed by the physician, and results of laboratory tests.

In our initial effort, we have reduced the number of features from the 119 encoded in the database to 40. We started by eliminating those features that had many missing values — numerical parameters expressing relevance of these features to the diagnosis would not be too reliable. Then we relied on expert's opinion as to which features have the highest diagnostic value. Having selected the total of 40 features, we elicited the structure of dependences among them from our domain experts: Dr. Hanna Wasyluk of the Medical Center of Postgraduate Education and Dr. Daniel Schwartz, a pathologist at the University of Pittsburgh. Subsequently we learned the parameters of the expert-constructed network, i.e., prior and conditional probabilities of the total of 41 variables, from the HEPAR database.

To evaluate the classification accuracy of our model, we performed a standard test in which we used a fraction of the database to learn the network parameters and the remainder of the records to test the network prediction. Our tests gave the following results. In over 36% of the cases, the most likely disorder was the correct diagnosis. In over 74% of the cases, the correct diagnosis was among the first four most likely disorders, as indicated by our model.

We would like to point out that our approach allows to query the system with partial observations, something that is not natural for classification systems. To test how intuitive and how useful the model is in practice, we have built a simple user interface to our model that lists all observable features, allows the user to set the values of any of them, updates the probability distribution over different disorders and presents an ordered list of possible diagnoses with their probabilities. Our program has been welcomed as a useful interactive diagnostic and training tool by our colleague physicians.

References

- [1] Leon Bobrowski. HEPAR: Computer system for diagnosis support and data analysis. Prace IBIB 31, Institute of Biocybernetics and Biomedical Engineering, Polish Academy of Sciences, Warsaw, Poland, 1992.
- [2] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1988.
- [3] Joe Whittaker. *Graphical Models in Applied Multivariate Statistics*. John Wiley & Sons, Chichester, 1990.