IctNeo system for jaundice management

Article · January 1998

Source: OAI

CITATIONS
6

READS
70

7 authors, including:

Concha Bielza
Universidad Politécnica de Madrid
369 PUBLICATIONS **6,412** CITATIONS
SEE PROFILE

Manuel Gómez-Olmedo
University of Granada
71 PUBLICATIONS **573** CITATIONS
SEE PROFILE

Juan A. Fernández del Pozo
Universidad Politécnica de Madrid
24 PUBLICATIONS **175** CITATIONS
SEE PROFILE

Sylvia Caballero
Hospital General Universitario Gregorio Marañón
27 PUBLICATIONS **274** CITATIONS
SEE PROFILE

Rev.R.Acad.Cienc.Exact.Fis.Nat. (Esp)
Vol. 92, N.º 4, pp 307-315, 1998
Monográfico: Problemas complejos de decisión

IctNeo SYSTEM FOR JAUNDICE MANAGEMENT

(decision analysis/influence diagrams/neonatal jaundice/intelligent/decision systems)

C. BIELZA*, M. GÓMEZ*, S. RÍOS-INSUA*, J.A. FDEZ. DEL POZO*, P. GARCÍA BARRENO**, S. CABALLERO***, M. SÁNCHEZ LUNA***

* Decision Analysis Group, Madrid Technical University, Spain.
** Spanish Royal Academy, Madrid, Spain.
*** Neonatology Service, Gregorio Marañón General Hospital, Madrid, Spain.

ABSTRACT

IctNeo is a complex decision support system for neonatal jaundice management, a very common medical problem. This paper discusses how it has been constructed in the context of Decision Analysis and how it operates. It isolates the difficulties encountered in practice related to knowledge-acquisition and related to computational limitations. Basically, a deeper insight into the problem solves the former while the latter is carried out with evidence propagation operations in influence diagrams

RESUMEN

IctNeo es un sistema complejo de ayuda a la decisión para la gestión de la ictericia neonatal, un problema médico muy común. Este artículo explica cómo se ha construido en el contexto del Análisis de Decisiones y cómo funciona. Se identifican las dificultades encontradas en la práctica respecto a la adquisición de conocimiento y a las limitaciones computacionales. Básicamente, una profundización en el problema resuelve el primer aspecto mientras que el segundo se soluciona con operaciones de propagación de evidencia en diagramas de influencia.

1. THE JAUNDICE PROBLEM

A few hours after birth, a baby skin and eyes may take a yellowish cast. This quite common condition is called jaundice. It is caused by the breakdown of excess red blood cells in the baby system which produces a substance called bilirubin. In the first days of a baby life, the liver is not mature enough to assimilate all the waste products produced and does not excrete bilirubin into the intestines at a normal rate. This excess bilirubin builds up in the baby bloodstream and tissues, giving the skin the characteristic yellow color.

Small or moderate increases of bilirubin are not harmful although the baby always needs to be watched closely for the first few days of life. Extremely high levels or hyperbilirubinemia can be harmful, causing potentially serious central nervous system and brain damages. This leads to the challenge of distinguising between the so-called physiologic jaundice and the more serious version, pathologic jaundice (4), related to the development of kernicterus (bilirubin encephalopathy) and the baby having risk factors.

A number of pathologies may have influence on hyperbilirubinemia: sepsis, congenital erythrocyte defect, inborn errors of metabolism, concealed hemorrhage, hypothyroidism, polycithemia, perinatal asphyxia, and isoimmunization, among others. The last one is associated to situations where mother and baby have different blood types and mother produces antibodies which destroy the infant red blood cells.

If the level of bilirubin is high enough to need treatment, it is usually treated with phototherapy. This means the undressed baby is placed under special lights or on a light producing blanket, both helping to break down the bilirubin. This intensive exposure involves an extended hospital stay and separation of the infant from the mother. If a baby bilirubin gets close to harmful levels, the doctor can do an exchange transfusion, a risky procedure where the baby blood is totally replaced. Current guidelines try to balance the risks of undertreatment and overtreatment (10), but there is a lack of consensus on when it is best to begin each treatment; at which point bilirubin levels are high enough to require treatment.

Neonatology Service of Gregorio Marañón Hospital in Madrid was interested in studying this problem as a Decision Analysis (DA) based problem. The hospital hopes to rely on an automated solution tool of this decision problem as an aid to the improvement of jaundice management. Doctors further pursue objectives like to include a big amount of uncertain factors and decisions —a hard task in the clinical practice—, to define better the moments to require and/or change the treatment, to decrease the costs of diagnostic and therapeutic phases, to decrease risks due to, e.g., the blood exchange, and to take into account the preferences of parents and doctors.

To that aim, the Decision Analysis Group of Madrid Technical University started to develop a decision support system called IctNeo, see (12) for its initial conception. It represents and solves the problem by means of an influence diagram (ID) (13), a more and more popular tool in DA. While conceptually simple, the application of IDs' methodology in practice may be extremely involved for large problems, encountering many difficulties that need a solution.

In the sections that follow, we shall develop the ideas mentioned above. Section 2 shows the process and some of the difficulties faced when constructing the system IctNeo. Section 3 explains how to propagate evidence on our problem in order to achieve the evaluation of complex IDs. Section 4 describes the system implementation, mainly its modules and user interface. Section 5 closes by outlining the conclusions and possible future research.

2. CONSTRUCTING IctNeo

2.1. General view of the problem

The first objective is to get a complete understanding of the problem, gathering its essential aspects. This is achieved after many interviews with doctors in which we obtain the following strategic knowledge that delimits the scope of the problem:

- The system only considers full-term infants born at this hospital and focuses on jaundice cases produced in the first moments of life (4 or less days old).
- The system evaluates closed cases, either because the baby is not already at hospital or he is receiving a treatment beyond jaundice. Thus, this is not a real time decision support system.
- The system does not only search for the best treatment, but also it aids with the first decision on admitting or not the newborn to hospital.
- The protocol of treatment focuses on the periodic control of bilirubin concentration to avoid it exceeds levels of potential neurological risks.
- The treatments considered are: exchange transfusion, phototherapy and observation.
- There exist a number of constraints on the way the possible treatments may combine, e.g., it is not possible to perform more than two exchanges on the same patient, we must start the treatment with observation or phototherapy, among others, see (12). We shall discuss below the important repercussions it brings about.

Once with this strategic knowledge, we analyze progressively the problem from minor to major detail.

The doctor first decides whether to admit or not the baby to hospital and confine it, eventually, to the Intensive Care Unit. In case of being admitted, it is necessary to control the bilirubin levels during this time, performing different tests and giving the patient some of the prescribed treatments, depending on some factors of the newborn and some concentration levels, regardless the cause of jaundice. Since jaundice extends in time for various days, the system must include a chain of decisions that comprises all the treatments the patients have received. The first and most generic diagram, see Figure 1, tries to capture the main idea of the problem: decision of admission followed by various treatment phases; uncertain factors to be revealed before or after making the decisions, like tests results; and a value node representing the utility function. The arcs explain our first beliefs on the influences among nodes.

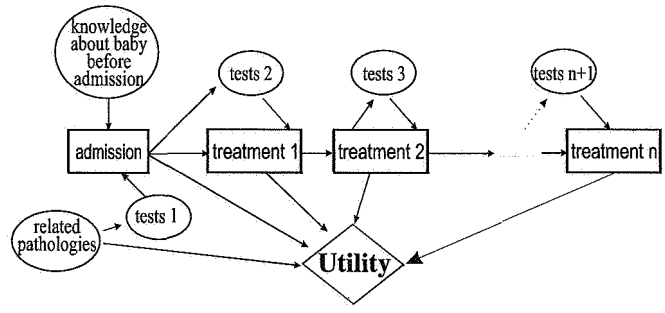

Figure 1. A generic ID for our problem

At each moment, the state of the patient determines the treatment to be applied. This leads at the same time to change his general state, reflected on clinical trials results and medical findings. A new treatment will follow the new state, and so on, until the infant is discharged or he receives a treatment that is outside our specific problem. Note that at each stage the state of the patient does not depend on the same set of factors, e.g, doctors take into account different aspects when deciding the admission and when deciding what to do with a patient already admitted.

This first scheme leads us to think of each action as the application of only one treatment. However, since the shortest treatment lasts 6 hours, the diagram would need a chain of 16 treatments of 6 hours to meet at most 4 days, and it entails an intractable ID due to the considerable set of nodes and arcs. We explain below how we face it by redefining those actions.

2.2. Specifying details

We now study in depth the variables of the problem. At the time of making the first decision, the doctor knows some factors of the neonate (and his mother) that inform about his state. Those factors can be grouped into 5 categories: mother data, mother tests, newborn data, newborn tests, and delivery data.

This general view of the first decision needs further specification of details that will yield the nodes associated to the left-hand side of the final diagram (Figure 2): *Mother data:* (1) administrative data, like age, race, whether she is primipara, whether she is ill; (2) social data like her cultural level and place of residence, that determine whether some slight cases could be discharged, all of them included in a node called "social cost"; (3) emotional data to value the emotional impact because of the interruption of parent-infant bonding. *Mother tests:* to know her blood group, Rh factor, and whether there exists isoimmunization. *Newborn data:* birth weight, gestational age, and age in hours. *Newborn tests:* Coombs' tests to determine isoimmunization of blood group and of Rh factor that reveal incompatibility between mother and baby bloods, tests to get bilirubin and hemoglobin concentrations, Apgar test and cord Ph test to study potential baby asphyxia, observation tests to estimate the degree of the yellowish skin color, and test to evaluate the risk of being admitted. *Delivery data:* whether it had instruments or not and whether it was necessary to give the baby some kind of revival.

As far as decisions are concerned, we have mentioned above the problem of having so many nodes if we consider periods of 6 hours. But we observed that treatments were referred very often as a combination of the initial therapies of 6 hours, leading to actions of different length, e.g. a phototherapy of 6, 12, or 24 hours long. Therefore, it is better to take these combinations in the domain of the decisions, deciding what to do in periods of time longer than the next 6 hours, thereby obtaining fewer decision nodes. Also, we overcome the problem of their constraints: a tree that represents these constraints on the whole process of 16 nodes would have 131071 endpoints, and there is no way to represent this kind of knowledge in the ID if we want to keep all the identical decisions of 6 hours mentioned earlier. However, the new alternatives are defined meeting the constraints avoiding incoherent optimal policies.

The final result of this process is the definition of 5 stages of treatment. There is one decision at the initial phase in which we choose the admission and alternatives allowed when starting the treatment, like phototherapies of different lengths. The main part of the treatment consists of 3 nodes (decisions 2, 3 and 4), the first two having in their domains the alternatives shown in Table 1. The alternatives for decision 4 are the first seven of Table 1. At the final phase, decision 5 includes only the first three possibilities of the same table. These more complex actions provide greater flexibility but require more complicated modeling.

1. dummy
2. observation + discharge
3. observation + outside treatment
4. observation (6 hours)
5. phototherapy (6 hours)
6. phototherapy (12 hours)
7. phototherapy (24 hours)
8. photot. (6 hours) + exchange + photot. (6 hours)
9. photot. (6 hours) + exchange + photot. (12 hours)

Table 1. Alternatives for decisions 2 and 3.

Note that because of phototherapies with different lengths, all the decisions have also different lengths having at most a total period of 90 hours, 2 or 2.5 days being the typical total length. There are patients who are not admitted or are discharged showing once more the varying length of the full process of jaundice, and, for that reason, the dummy alternative (not to do anything) is included to fill in the remaining domains. It leads to a highly asymmetric ID with incoherent sequences of decisions (e.g., admission for decision 1 and dummy for decision 2). But we shall give the system the capacity for somehow dealing with these combinations of therapies representing impossible scenarios.

We focus now on the following phases of the protocol. The action made at time t influences on the patient state at $t+1$, determined by means of a number of clinical trials. The details are: (1) *Patient tests:* to determine bilirubin and hemoglobin levels showing the evolution of the patient; (2) *Patient data:* mainly his age which is closely linked to bilirubin increase. This information is shown in the area surrounded by a dotted line in the right-hand side of Figure 2. The addition of the dependencies among the different variables shown by the arcs of the diagram, its continuous revisions as long as we obtain new and refined knowledge, have led us to have that "final" ID shown in Figure 2.

2.3. KNOWLEDGE-ACQUISITION

2.3.1. Domains

The domains of decision variables are already identified and the next step is to define those of random variables. Some of them are referred to concentration measures, weights, ages, etc. and are easily defined, even being continuous because doctors decided to work with discrete approximations, see (12). Other variables state medical appraisals of the situation. For these, we have tried to be close to doctors thinking, e.g., to define the node "pathology seriousness", doctors gave a set of levels of seriousness in a way close to their usual language —slight, serious, and very serious— and with enough detail to cover all the situations.

2.3.2. Probabilities

The task of eliciting the probabilistic information for the chance nodes is more difficult. Typically, we used both historical data —taken from the Neonatology Service of the hospital— and subjective judgments, and followed the SRI protocol and its extensions (9). Probability encoding is extremely involved when the problem is structurally complicated, say a heavily asymmetric and dense large ID, with continuous random variables. The discretization suggested by doctors relieved this last problem. However, a chance node with many predecessors poses problems related to how

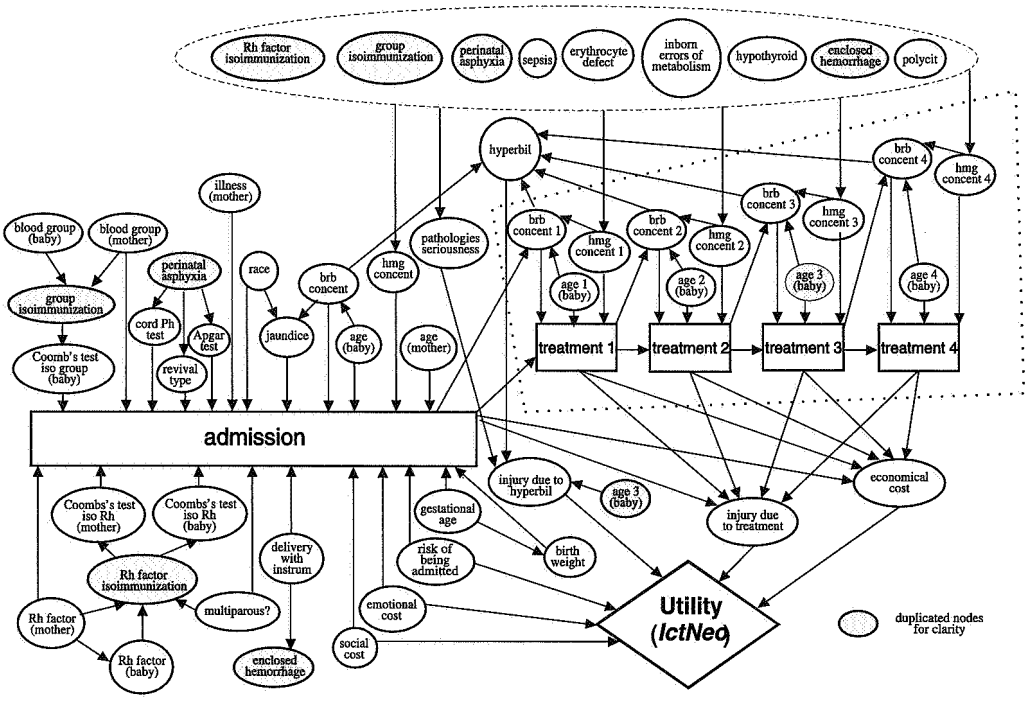

Figure 2. Final ID for our problem.

to obtain from experts, tables with so many entries and how to store and manage so much information.

For example, consider the node "injuries due to hyperbilirubinemia" (IHY), see Figure 2. The initial versions of the ID did not contain the node "hyperbilirubinemia", and the predecessors of node IHY were: all the bilirubin concentrations, "pathology seriousness", and the newborn age at the third stage (when a higher level of bilirubin may exist), see Figure 3 a). Therefore, we should assign $6 \cdot 3 \cdot 2^6 = 1152$ numbers to obtain the probability distribution for node IHY.

To face this problem, we used generalized noisy OR-gates (2, 11), which are based on a model of causal nature, with some causes $P_1,..., P_k$ acting to produce effect X, all the causes and the effect having values absent and present with various grades of intensity. After having checked their assumptions, the only assignments required to derive the others are those of the conditional probabilities of X given that all causes but one are absent, needing a number of assignments that is linear in the number of causes instead of exponential, as in the general case. In our case, by adding node "hyperbilirubinemia", we can define a causal relation between bilirubin levels and the presence or not of their alteration. Since with the other variables —"pathology seriousness" and the newborn age at the third stage—, is not possible this causal relation, the diagram is the one shown in Figure 3 b). Now, we need only to assign 10 values for node "hyperbilirubinemia" plus $6 \cdot 3 \cdot 2^2 = 72$ for node IHY, for a total of 82 values.

Our medical problem had 50 chance nodes, obtaining, with OR-gate modelisations, a reduction of 70% in the number of assignments required.

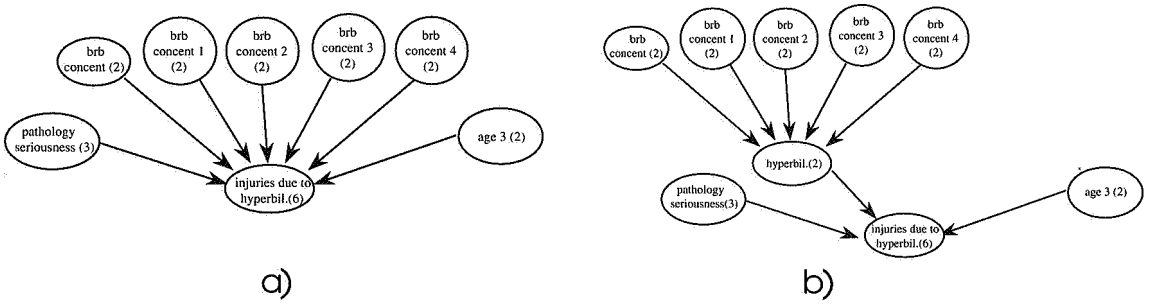

Figure 3. Node *injuries due to hyperbilirubinemia* a) before and b) after OR-gate modelisations. Number in parentheses indicates the cardinal of the domain.

2.3.3. Utilities

The only issue left to complete the ID is the preferences elicitation. We have shown above the value node in that ID (see Figure 2) representing the (expected) utilities depending on their predecessors. Some of these predecessors have been explained in prior sections, but the process to get all of them follows, as typically, a structured scheme: (1) to construct an objectives hierarchy along with the corresponding attributes and their scales; and (2) to derive the functional form of the multi-attribute utility function consistent with various sets of independence assumptions about the risk attitudes investigated (6). This function collects the influence of each subobjective on the overall objective.

As a result, we obtained the hierarchy shown in Figure 4 with 6 lowest-level objectives. Highlight X_4, concerning the stay at hospital which arises from the risk of infections, contagions, etc. and the preferences of parents measured by the inconvenience of visiting the baby every day (X_2) if he is at hospital, and the interruption of parent-infant bonding (X_3). For all attributes except for X_1, *ad hoc* constructed scales were introduced. Observe now how all X_i's are represented in the ID (Figure 2) as factors that value the process, illustrating that not only therapeutic actions are important but also their suitability to the patient situation.

The utility function was derived as multiplicative on (X_1, $Y_1 = (X_2, X_3)$, X_4, $Y_2 = (X_5, X_6)$) with additive decompositions for (X_2, X_3) and (X_5, X_6). The assignment of component utility functions and scaling constants was carried out with the aid of Logical Decisions (8), but allowing imprecise assignments, as a way of sensitivity analysis, see (5) for a detailed explanation.

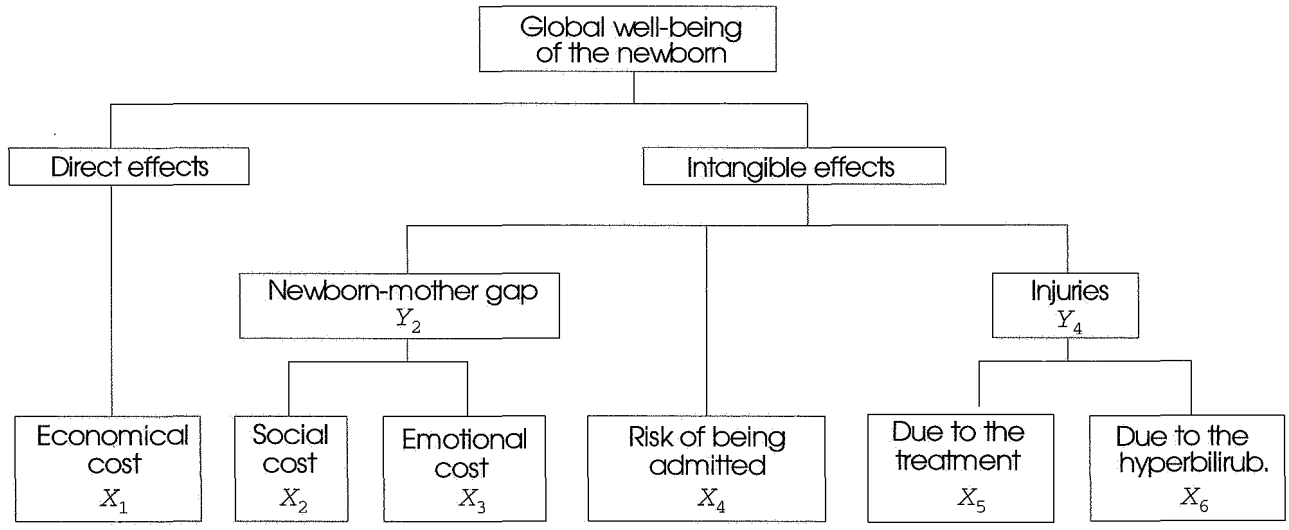

Figure 4. An objectives hierarchy for our problem

3. EVIDENCE PROPAGATION

The evaluation of an ID is typically carried out via Shachter's algorithm (13), but it works for relatively small problems. As we have seen, our jaundice problem has a big size: 50 chance nodes, 5 decision nodes and 144 arcs connecting some of them, plus 112 no-forgetting arcs. It amounts to requiring 95858 memory positions to storage the probability distributions (90458) plus the utility function (5400).

Typically, the need for storage space grows enormously during the solution process, due to node inheritances at chance node removal and arc reversal operations. For instance, when removing a chance node C which precedes the value node v, v inherits the predecessors of C (if any) becoming new predecessors of v, provided that they were not before. Since the storage requirements for the utility function grow exponentially in the number of variables that influence on it, this situation can lead to require more memory space than that available in any personal computer.

A first way to overcome this problem is to find the optimal sequence of variable reductions and arc removals, that which involves the minimum computational effort. But this is a *NP-hard* problem and, therefore, we can only pick a *good* sequence via some heuristic, like Kong's greedy heuristic (7). We have worked on this, see (1), improving the existent heuristics, obtaining e.g., a required storage space divided at most by two with respect to Kong's heuristic. Also, we have provided other computational savings by computing the expected utilities when a chance node is removed, only whenever this implies a saving in the storage space with respect to the previous diagram. Otherwise, we postpone it until it is necessary, i.e., at the moment of a decision node removal.

But this is not enough for big problems. Our ID requires a maximum storage capacity (for the operation which introduces the highest increase) of 1.66×10^{14} storage positions, being the average size of the problem 1.79×10^{13} storage positions. Therefore, the evaluation process becomes unmanageable.

As we said above, the system manages closed cases, from which there is some information available, i.e., the outcomes of most random variables are known (maybe except for some hemoglobin measures not taken). This knowledge can be used via the technique of evidence propagation (3). The ID is updated based on these new observations reflecting then the new state of knowledge by changing the probability of observed variables together with the ID structure.

Consider, e.g., the operation of chance-to-chance evidence propagation. Figure 5 will illustrate the ideas. Suppose $X = x_1$ is known. We will now add this information reducing the uncertainty content on the problem.

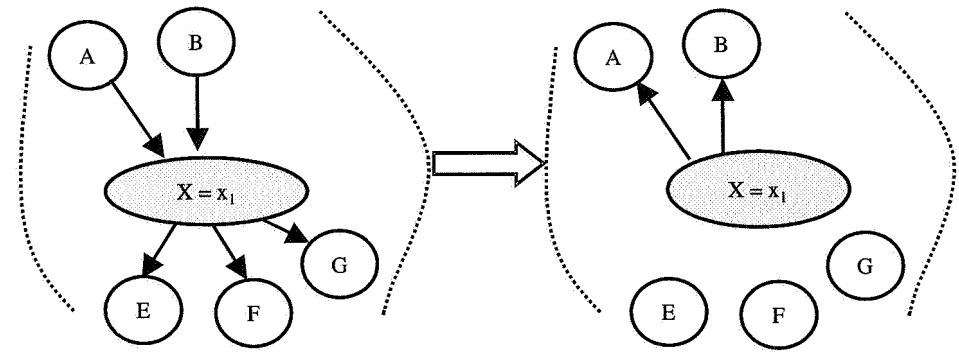

P(E|X)

	x_1	x_2	x_3
e_1	0.2	0.1	0.7
e_2	0.8	0.9	0.3

P(A|X)

	x_1	x_2	x_3
a_1	0.1	0.2	0.8
a_2	0.9	0.2	0.1
a_3	0	0.6	0.1

Figure 5. Chance-to-chance evidence propagation.

The evidence x_1 is first absorbed on X just by the table lookup of the observed outcome. This has effect on successors and predecessors of X, as follows. First, a forward evidence propagation to its successors E, F, G takes place, which is to discard situations of X taking values different than the observed x_1, reducing in turn the size of the probability distributions of these three variables. For example, the distribution for E changes from $P(E|X)$ to $P(E|X = x_1)$, as shown in Figure 5 by the non-shaded zone of the table, deleting then the arc (X,E). Second, a backward evidence propagation to its predecessors A, B takes place. For that, we first reverse the arcs between X and its predecessors by application of Bayes' rule, these becoming successors of X as in the previous case. Therefore, arcs (X, A) and (X, B) are eliminated and X too. The computational burden alleviation is obvious since we reduce the sizes of probability tables and we eliminate node X and the arcs from and onto X.

There are similar operations to propagate evidence with other kind of nodes. Should the value node be a successor of an observed node, the storage requirements for the utility function would be reduced. These ideas have permitted us to solve our jaundice problem. For each patient, we have a historical record with his data, and those of his mother and familiar environment. This evidence is known before making the first decision and do not take into account any medical valuation, thus keeping the uncertainty on all the aspects related to the medical valuation of the problem.

As a result, it produces important implications on the system operation. The conventional evaluation of an ID provides at each decision node, a table with the optimal alternative for *each possible combination* of outcomes of the value node predecessors at that moment. Once the system has generated all these tables, we only must look for that combination that coincides with the patient under consideration. Hence, the ID is evaluated only once and we perform on these tables as many lookups as cases studied.

Nevertheless, the evaluation of an ID by instantiation of evidence on some nodes amounts to solving it for *each particular patient* once this evidence has been propagated and has updated the ID. Consequently, we must evaluate as many IDs as different cases studied. Decision tables will have a smaller size because they do not contain any reference to the instantiated variables. Our requirements have decreased now with this approach: 88652 storage positions, with a maximum storage capacity and an average size of 5.35×10^9 and 7.48×10^8 storage positions, respectively.

4. IMPLEMENTATION

IctNeo is implemented in C++ under Windows-95. Figure 6 shows IctNeo general architecture.

IctNeo includes its own grammar to describe the ID, th-

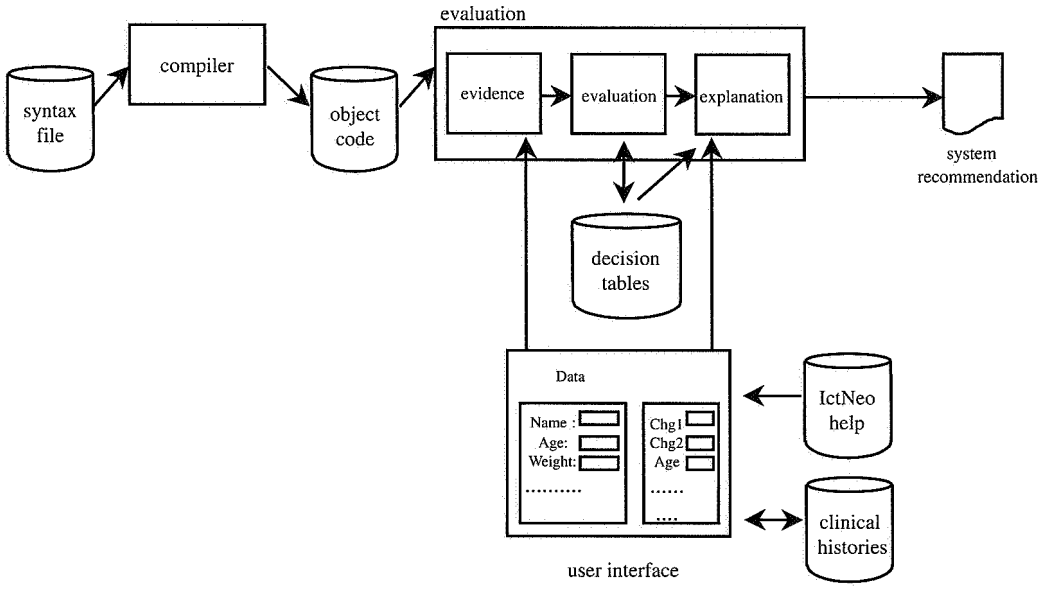

Figure 6. The architecture of IctNeo system.

rough which all the nodes, their relationships and the relevant information are defined and written in the *syntax file*, being easily modifiable. A chance node, e.g., may be declared by providing either its whole probability table or the OR-gate parameters that generate internally this table. The *compiler* reads this file by following the grammar and checks for syntactic and lexical errors. In case it does not exist any error, the system yields an *object code* from the information gathered from the specification, the data structure needed to support the diagram information. Then it is verified whether the diagram satisfies the necessary conditions to be evaluated (no cycles, coherent probability distributions, etc.). Both the grammar definition and the compiler have been thought so as to develop a general framework ready to solve every ID, and provide an easy interchange of files among different working groups.

The evaluation module implements Shachter's algorithm (13) and obtains the optimal strategies. First, the patient data at the moment of the admission are entered and the ID is updated by this *evidence*. Second, the *evaluation* is carried out on that simplified diagram, yielding the *decision tables* corresponding to our patient. IctNeo incorporates the knowledge about the asymmetries of our problem due to constraints on decisions: at the time the first decision node is removed during the process, the value node has already inherited all the decision nodes and its table is cut by discarding the sequences that do not meet the constraints. The system knows it because it is written in the grammar script. Third, we can compare the *system recommendation* with the action already made by doctors to our patient. To get an insight into it, the system generates *explanations* of the results looking for irrelevant variables, those that lead to the same decisions being made regardless the values they take.

The *user interface* allows for data entry of a patient already treated from which the evidence is taken. It has some *help* for all the menus and a *data base* with all the cases studied to account for the comparisons between the system and doctor decisions in order to draw conclusions. Figure 7 shows the main menu and some windows to update administrative and admission data.

5. CONCLUSIONS AND FURTHER RESEARCH

Once the difficulties encountered when modeling and solving the ID for this problem have been overcome, IctNeo is still being tested and refined by performing sensitivity analyses. It allows the user to check the impact of key parameters in the evolution of the patient, perhaps suggesting changes in utilities, probabilities, and even in the structure of the diagram.

Current research is directed towards extending the analysis as follows:

1. To weaken the uncertainty assignment process by allowing the expert to provide probability intervals rather than probability exact values.

2. To automate the sensitivity analysis tools and to endow the system with learning capabilities: the initial intervals assigned may be balanced according to the valuation of the expert about the system recommendations.

3. To analyze, by means of data bases techniques, the decision tables the system proposes. We can then

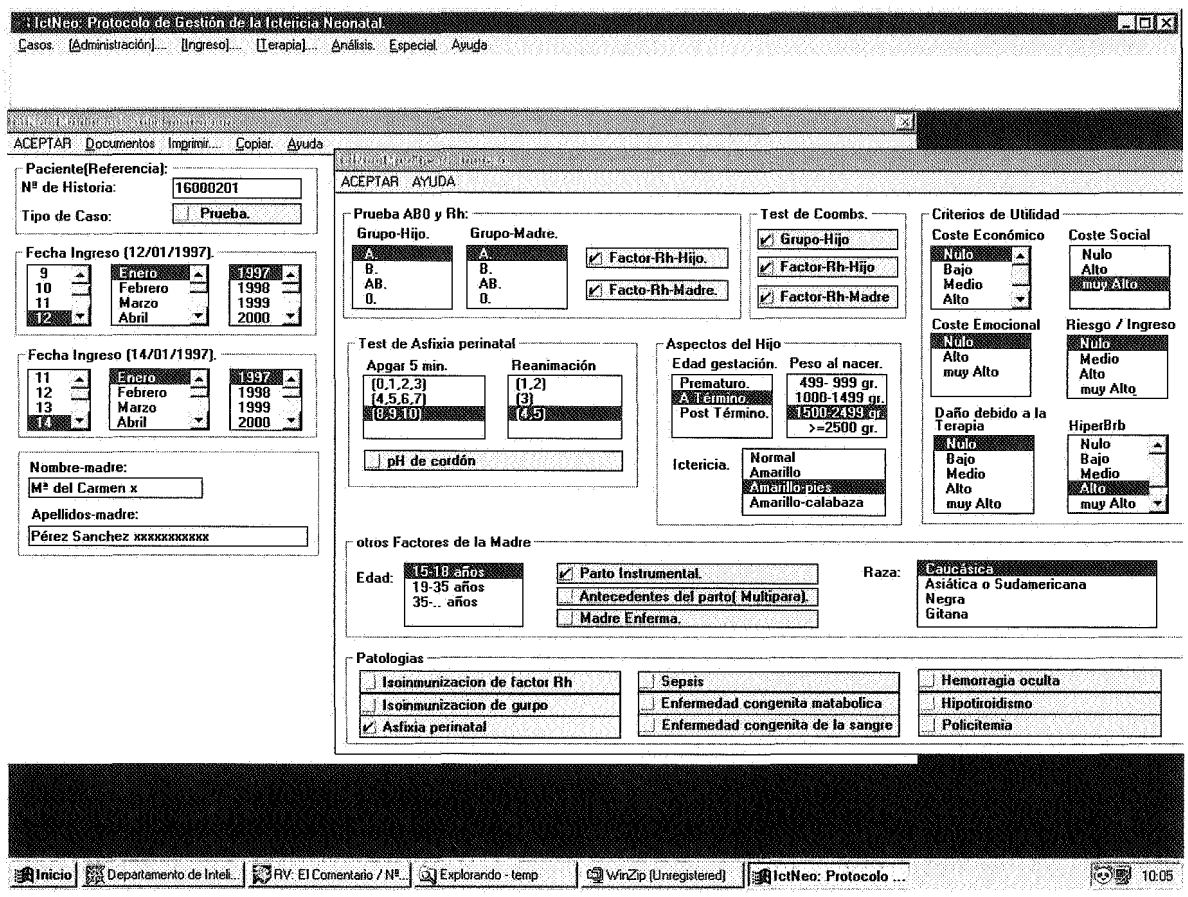

Figura 7. User interface of IctNeo.

find out which are the relevant factors to apply the different therapies. The matter in question is to draw conclusions from the whole set of the system proposals with regard to all possible user profiles.

4. To include pathology diagnosis by looking at the final updated probabilities of the diagram.

5. To have the possibility of dealing with missing patient information and/or unspecified distributions.

6. To include a constructor of (imprecise) utility functions to be used in other different hospitals.

In short, we hope IctNeo will become an important aid for doctors, both for interns and residents, who will be able to evaluate the implications relative to changes of protocol of neonatal jaundice in newborns, and will have in turn a better understanding of the jaundice problem.

ACKNOWLEDGMENTS

Research supported by DGESIC project PB97-0856 and FIS project 97/0003-02. We wish to thank doctor D. Blanco who helped us in our visits to hospital.

REFERENCES

1. Bielza, C., Gómez, M., Ríos-Insua, S. & Fdez del Pozo, J.A. (1999) Structural, Elicitation and Computational Issues Faced when Solving Complex Decision Making Problems with Influence Diagrams, *Computers & Operations Research*, to appear.

2. Díez, F.J. (1993) Parameter Adjustment in Bayes Networks. The Generalized Noisy OR-Gate, in D. Heckerman and A. Mamdani (eds.), *Uncert. in Artif. Intell.: Proc. of the 9th Confer.* 99-105, Morgan Kaufmann, San Mateo, CA.

3. Ezawa, K.J. (1998) Evidence Propagation and Value of Evidence on Influence Diagrams, *Oper. Res.*, **46,** 1, 73-83.

4. Gartner, L.M. (1995) Neonatal Jaundice, *Ped. in Review,* **16**, 22-31.

5. Gómez, M., Ríos-Insua, S., Bielza, C. & Fdez del Pozo, J.A. (1999) Multi-attribute Utility Analysis in the IctNeo System, in Y.Y. Haimes and R. Steuer (eds), *XIV-th Int. Conf. on Multiple Criteria Decision Making,* Lecture Notes in Economics and Mathematical Systems, Springer, Berlin (to appear).

6. Keeney, R.L. & Raiffa, H. (1976) *Decisions with Multiple Objectives. Preferences and Value Tradeoffs,* Wiley.

7. Kong, A. (1986) Multivariate Belief Functions and Graphical Models, PhD. Diss., Dept. of Statistics, Harvard Univ., Cambridge, Mass.

8. Logical Decisions for Windows 95 (1998) Multimeasure Decision Analysis Software, Golden, CO.

9. Merkhofer, M.W. (1987) Quantifying Judgmental Uncertainty: Methodology, Experiences, and Insights, *IEEE Trans. on Sys., Man and Cyber.* SMC **17,** 741-752.

10. Newman, T.B. & Maisels, M.J. (1992) Evaluation and Treatment of Jaundice in the Term Infant: A Kinder, Gentler Approach, *Pediatrics,* **89,** 809-830 (with discussion).

11. Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems,* Morgan Kaufmann.

12. Ríos-Insua, S., Bielza, C., Gómez, M., Fdez del Pozo, J.A., Sánchez, M. & Caballero, S. (1998) An Intelligent Decision System for Jaundice Management in Newborn Babies, in F.J. Girón (ed.), *Applied Decision Analysis,* 133-144, Kluwer.

13. Shachter, R.D. (1986) Evaluating Influence Diagrams, *Oper. Res.,* **34,** 871-882.