A Case Study on the Application of Probabilistic Conditional Modelling and Reasoning to Clinical Patient Data in Neurosurgery

Christoph Beierle[1], Marc Finthammer[1], Nico Potyka[1], Julian Varghese[1], and Gabriele Kern-Isberner[2]

[1] Dept. of Computer Science, FernUniversität in Hagen, 58084 Hagen, Germany
[2] Dept. of Computer Science, TU Dortmund, 44221 Dortmund, Germany

Abstract. We present a case-study of applying probabilistic logic to the analysis of clinical patient data in neurosurgery. Probabilistic conditionals are used to build a knowledge base for modelling and representing clinical brain tumor data and expert knowledge of physicians working in this area. The semantics of a knowledge base consisting of probabilistic conditionals is defined by employing the principle of maximum entropy that chooses among those probability distributions satisfying all conditionals the one that is as unbiased as possible. For computing the maximum entropy distribution we use the MEcoRe system that additionally provides a series of knowledge management operations like revising, updating and querying a knowledge base. The use of the obtained knowledge base is illustrated by using MEcoRe's knowledge management operations.

1 Introduction

In the medical domain, uncertain rules like "*If symptoms S_1, S_2, and S_3 are present, then there is a probability of 70% that the patient has disease D.*" occur frequently. An intelligent agent providing decision support for performing medical diagnosis and for choosing a therapy must be able to deal with pieces of knowledge expressed by such rules, requiring elaborate knowledge representation and reasoning facilities. For instance, in neurosurgery, such an agent should be able to answer diagnostic questions in the presence of evidential facts like "*Given the evidence that the patient has perceptual disturbances, suffers from unusual pain in the head and that there are symptoms for intracranial pressure, what is the probability that he has a cranialnerve tumor?*", and the agent should be able to perform hypothetical reasoning as in: "*There is evidence that the patient has perceptual disturbances and that there are symptoms for intracranial pressure. If we chose a surgery for therapy and if the correct diagnosis was gliobastoma, what would be the patient's chance to recover completely without any serious complications?*" Moreover, when the agent lives in an uncertain and dynamic environment, she has to adapt her epistemic state constantly to changes in the surrounding world and to react adequately to new demands (cf. [5], [10]).

L.C. van der Gaag (Ed.): ECSQARU 2013, LNAI 7958, pp. 49–60, 2013.

In this paper, we report on a case study on the application of probabilistic modelling and reasoning to clinical patient data in neurosurgery. A knowledge base `BT` representing and integrating both statistical frequencies of brain tumors reported in the literature as well as physicians' expert beliefs is developed and used to perform reasoning regarding the diagnosis of brain tumor types or the prognosis for patients (see [22,21] for more information on the medical background). Uncertain rules as the first one above are modelled by probabilistic conditionals, formally denoted by $(D|S_1 \wedge S_2 \wedge S_3)[0.7]$. Semantics of such conditionals are given by probability distributions over the possible worlds determined by the underlying propositional variables, and satisfaction of a conditional by a probability distribution P is defined via conditional probability, e.g., P satisfies $(D|S_1 \wedge S_2 \wedge S_3)[0.7]$ iff $P(D|S_1 \wedge S_2 \wedge S_3) = 0.7$. In order to complete any missing or unspecified knowledge, the concept of maximum entropy [16,11] is used. The required reasoning is carried out by the MEcoRe system [7] that implements reasoning at optimum entropy and provides knowledge management operations required for modelling an intelligent agent.

In the following section, we first recall some preliminaries of probabilistic conditional logic and features of the MEcoRe system as they are presented in [7]. In Sec. 3, the vocabulary of `BT` and a first version of this knowledge base is presented. Section 4 introduces revision and update operations for `BT`, and in Sec. 5 we illustrate the reasoning facilities for prognosis and hypothetical what-if-analysis, demonstrating that the results are well in accordance with a clinical physician's point of view. In Sec. 6, we conclude and point out further work.

2 Background: Probabilistic Conditionals and MEcore

2.1 Probabilistic Conditional Logic in a Nutshell

We start with a propositional language $\mathcal{L}$, generated by a finite set Σ of (binary) atoms $a, b, c, \ldots$. The formulas of $\mathcal{L}$ will be denoted by uppercase Roman letters $A, B, C, \ldots$. For conciseness of notation, we will omit the logical *and*-connector, writing AB instead of $A \wedge B$, and over-lining formulas will indicate negation, i.e. $\overline{A}$ means $\neg A$. Let Ω denote the set of possible worlds over $\mathcal{L}$; Ω will be taken here simply as the set of all propositional interpretations over $\mathcal{L}$ and can be identified with the set of all complete conjunctions over Σ. For $\omega \in \Omega$, $\omega \models A$ means that the propositional formula $A \in \mathcal{L}$ holds in the possible world ω.

By introducing a new binary operator $|$, we obtain the set $(\mathcal{L} \mid \mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$ of (unquantified) *conditionals* (or *rules*) over $\mathcal{L}$. $(B|A)$ formalizes "*if A then B*" and establishes a plausible, probable, possible etc connection between the *antecedent* A and the *consequent* B. We will use $Sen_{\mathcal{C}}$ to denote the set of all *probabilistic conditionals* (or *probabilistic rules*) of the form $(B|A)[x]$ where x is a probability value $x \in [0, 1]$.

To give appropriate semantics to conditionals, they are usually considered within richer structures such as *epistemic states*. Besides certain (logical) knowledge, epistemic states also allow the representation of e.g. preferences, beliefs, assumptions of an intelligent agent. Basically, an epistemic state allows one to compare formulas or worlds with respect to plausibility, possibility, necessity, probability etc. In a quantitative framework, most appreciated representations of epistemic states are provided by *probability functions* (or *probability distributions*) $P : \Omega \to [0,1]$ with $\sum_{\omega \in \Omega} P(\omega) = 1$. Thus, in this setting, the set of *epistemic states* we will consider is $\mathit{EpState} = \{P \mid P : \Omega \to [0,1]$ is a probability function$\}$. The probability of a formula $A \in \mathcal{L}$ is given by $P(A) = \sum_{\omega \models A} P(\omega)$, and the probability of a conditional $(B|A) \in (\mathcal{L} \mid \mathcal{L})$ with $P(A) > 0$ is defined as $P(B|A) = P(AB)/P(A)$, the corresponding conditional probability. Conditionals are interpreted via conditional probability. So the satisfaction relation $\models_{\mathcal{C}} \subseteq \mathit{EpState} \times \mathit{Sen}_{\mathcal{C}}$ of probabilistic conditional logic is defined by $P \models_{\mathcal{C}} (B|A)\,[x]$ iff $P(B|A) = x$.

2.2 Epistemic States and Belief Management Operations

Initialization. First, a prior epistemic state has to be built up on the basis of which the agent can start her computations. If no knowledge at all is at hand, simply the uniform epistemic state is taken to initialize the system. In our probabilistic setting, this corresponds to the uniform distribution where each possible world is assigned the same probability. If, however, a set of probabilistic rules is at hand to describe the problem area under consideration, an epistemic state has to be found to appropriately represent this prior knowledge. To this end, we assume an inductive representation method to establish the desired connection between sets of sentences and epistemic states. Whereas generally, a set R of sentences allows a (possibly large) set of models (or epistemic states), in an inductive formalism we have a function *inductive* : $\mathcal{P}(\mathit{Sen}_{\mathcal{C}}) \to \mathit{EpState}$ (where $\mathcal{P}(S)$ denotes the power set of S) such that *inductive*(R) selects a unique, "best" epistemic state from all those states satisfying R.

In the probabilistic framework, the *principle of maximum entropy* associates to a set R of probabilistic conditionals the unique distribution $P^* = \mathit{MaxEnt}(R)$ that satisfies all conditionals in R and has maximal entropy, i.e., $\mathit{MaxEnt}(R)$ is the *unique* solution to the maximization problem

$$\arg\max_{P' \models R} H(P') \quad \text{with } H(P') = -\sum_{\omega} P'(\omega) \log P'(\omega) \tag{1}$$

The rationale behind this is that $\mathit{MaxEnt}(R)$ represents the knowledge given by R most faithfully, i.e. without adding information unnecessarily (cf. [16,11]).

Example 1. Consider the three propositional variables s - *being a student*, y - *being young*, and u - *being unmarried*. *Students* and *unmarried people* are mostly *young*. This commonsense knowledge an agent may have can be expressed probabilistically e.g. by the set $R = \{(y|s)[0.8], (y|u)[0.7]\}$ of conditionals. The *MaxEnt*-representation $P^* = \mathit{MaxEnt}(R)$ computed by MEcoRe is:

ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$
syu	0.1950	$sy\overline{u}$	0.1758	$s\overline{y}u$	0.0408	$s\overline{y}\,\overline{u}$	0.0519
$\overline{s}yu$	0.1528	$\overline{s}y\overline{u}$	0.1378	$\overline{s}\,\overline{y}u$	0.1081	$\overline{s}\,\overline{y}\,\overline{u}$	0.1378

Querying an Epistemic State. Querying an agent about her beliefs amounts to pose a set of unquantified sentences and asking for the corresponding degrees of belief with respect to her current epistemic state.

Example 2. Suppose the current epistemic state is $currState = MaxEnt(R)$ from Ex. 1, and our question is "*What is the probability that unmarried students are young?*", i.e. the set of queries is $\{(y|su)\}$. MECoRe returns $\{(y|su)[0.8270]\}$, that is, unmarried students are supposed to be young with probability 0.8270.

New Information and Belief Change. Belief revision, the theory of dynamics of knowledge, has been mainly concerned with propositional beliefs for a long time. The most basic approach here is the *AGM-theory* presented in the seminal paper [1] as a set of postulates outlining appropriate revision mechanisms in a propositional logical environment. This framework has been widened by Darwiche and Pearl [5] for (qualitative) epistemic states and conditional beliefs. An even more general approach, unifying revision methods for quantitative and qualitative representations of epistemic states, is described in [12]. The crucial meaning of conditionals as *revision policies* for belief revision processes is made clear by the so-called *Ramsey test*, according to which a conditional $(B|A)$ is accepted in an epistemic state Ψ, iff revising Ψ by A yields belief in B: $\Psi \models (B|A)$ iff $\Psi * A \models B$ where $*$ is a belief revision operator (see e.g. [8]).

Note, that the term "belief revision" is a bit ambiguous: On the one hand, it is used to denote quite generally *any* process of changing beliefs due to incoming new information [8]. On a more sophisticated level, however, one distinguishes between different kinds of belief change. Here, *(genuine) revision* takes place when new information about a static world arrives, whereas *updating* tries to incorporate new information about a (possibly) evolving, changing world [10]. Further belief change operators are *expansion*, *focusing*, *contraction*, and *erasure* (cf. [8,6,10]). In the following, we will use the general approach to belief change developed in [12] where belief change is considered in a very general and advanced form: Epistemic states are revised by sets of conditionals – this exceeds the classical AGM-theory by far which only deals with sets of propositional beliefs.

In the probabilistic framework, a powerful operator to change probability distributions by sets of probabilistic conditionals is provided by the *principle of minimum cross-entropy* which generalizes the principle of maximum entropy in the sense of (1): Given a (prior) distribution P and a set R of probabilistic conditionals, the *MinCEnt-distribution* $P^* = MinCEnt(P, R)$ is the *unique* distribution that satisfies all constraints in R and has minimal cross-entropy H_{ce} with respect to P, i.e. P^* solves the minimization problem

$$\arg\min_{P' \models R} H_{ce}(P', P) \quad \text{with } H_{ce}(P', P) = \sum_{\omega} P'(\omega) \log \frac{P'(\omega)}{P(\omega)} \tag{2}$$

If R is basically compatible with P (i.e. *P-consistent*, cf. [12]), then P^* is guaranteed to exist (for further information and lots of examples, see [4,16,12]). The cross-entropy between two distributions can be taken as a directed (i.e. asymmetric) information distance [19] between these two distributions. Following the principle of minimum cross-entropy means to modify the prior epistemic state P in such a way as to obtain a new distribution P^* which satisfies all conditionals in R and is as close to P as possible. So, the *MinCEnt*-principle yields a probabilistic belief change operator, associating to each probability distribution P and each P-consistent set R of probabilistic conditionals a revised distribution $P^* = \mathit{MinCEnt}(P, R)$ in which R holds. In [13] it is shown how both revision and update can be based on such a belief change operator, and the corresponding conceptual agent model MEcoRe which realizes this approach is described in [2].

Example 3. Suppose that some time later, the relationships in the population from Example 1 between students and young people have changed, so that students are young with a probability of 0.9. In order to incorporate this new knowledge, the agent applies an updating operation to modify P^* appropriately. The result $P^{**} = \mathit{MinCEnt}(P^*, \{(y|s)[0.9]\})$ as determined by MEcoRe is:

ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$
syu	0.2151	$sy\overline{u}$	0.1939	$s\overline{y}u$	0.0200	$s\overline{y}\,\overline{u}$	0.0255
$\overline{s}yu$	0.1554	$\overline{s}y\overline{u}$	0.1401	$\overline{s}\,\overline{y}u$	0.1099	$\overline{s}\,\overline{y}\,\overline{u}$	0.1401

It is easily checked that indeed, $P^{**}(y|s) = 0.9$ (taking rounding into account).

Diagnosis. Diagnosing a given case is one of the most common operations in knowledge based systems. Given some case-specific *evidence* E (formally, a set of quantified facts), diagnosis assigns degrees of belief to the atomic propositions D to be *diagnosed* (formally, D is a set of unquantified atomic propositions). Thus, making a diagnosis in the light of some given evidence corresponds to determine what is believed in the state obtained by focusing the current state P on the given evidence, i.e. querying the epistemic state $\mathit{MinCEnt}(P, E)$ with respect to D. Thus, here focusing corresponds to conditioning P with respect to the given evidence E.

Example 4. Let $\mathit{currState} = P^*$ from Ex. 1. If there is now certain evidence for being a student and being unmarried – i.e. $E = \{su[1]\}$ – and we ask for the degree of belief of being young – i.e. $D = \{y\}$ –, MEcoRe computes $\{y[0.8270]\}$. Thus, if there is certain evidence for being an unmarried student, then the degree of belief for being young is 0.8270.

What-If-Analysis: Hypothetical Reasoning. Hypothetical reasoning asks for the degree of belief of complex relationships (goals) under some hypothetical assumptions. This is useful, e. g., to exploit in advance the benefits of some expensive or intricate medical investigations. Note that whereas in the diagnostic case both evidence E and diagnoses D are just simple propositions, in hypothetical reasoning both the *assumptions* A (formally, a set of quantified conditionals) as well as the *goals* G (formally, a set of unquantified conditionals) may be sets of full conditionals. However, since its underlying powerful *MinCEnt*-update operator can modify epistemic states by arbitrary sets of conditionals, MEcoRe can handle hypothetical what-if-analysis structurally analogously to the diagnostic case, i. e. by querying the epistemic state $\mathit{focussed\_state} = \mathit{MinCEnt}(P, A)$ with respect to G where P is the current epistemic state. Since this is hypothetical reasoning, the agent's current epistemic state remains unchanged.

Example 5. Given $currState = P^*$ from Ex. 1 as present epistemic state, a hypothetical reasoning question is given by: "What would be the probability of being young under the condition of being unmarried – i.e. $G = \{(y|u)\}$ –, provided that the probability of a student being young changed to 0.9 – i.e. $A = \{(y|s)[0.9]\}$?" MEcoRe's answer is $\{(y|u)[0.7404]\}$ which corresponds to the probability given by P^{**} from Ex. 3.

2.3 The MEcore System

The main objective of MEcoRe is to implement probabilistic reasoning at optimum entropy and to support advanced belief management operations like revision, update, diagnosis, or what-if-analysis in a most flexible and easily extendable way. MEcoRe is implemented in Java and uses a straight-forward, direct implementation of a well-known *MinCEnt* algorithm, computing the distribution $P^* = MinCEnt(P, R)$ in an iterative way [4], and provides a powerful and flexible interface. MEcoRe can be controlled by a text command interface or by scripts, i. e. text files that allow the batch processing of command sequences. These scripts and the text interface use a programming language-like syntax that allows to define, manipulate and display variables, propositions, rule sets and epistemic states. The following example shows a way to generate an epistemic state using the initialize and update operators:

```
//define a set of rules
kb := ((y|s)[0.8], (y|u)[0.7]);
// initialize an epistemic state with these rules
currState := epstate().initialze(kb);
//query and output current belief in the conditional (y|su)
currState.query((y|su));
//update the epistemic state currState by (y|s)[0.9]
currState.update((y|s)[0.9]);
```

Hence, one is able to use both previously defined rule sets and rules that are entered just when they are needed, and combinations of both. The ability to manipulate rule sets, to automate sequences of updates and revisions, and to output selected results for comparing, yields a very expressive command language. This command language is a powerful tool for experimenting and testing with different setups. All core functions of the MEcoRe system are also accessible through a software interface in terms of a Java API; thus, MEcoRe can easily be extended by a GUI or be integrated into another software application.

There are many systems performing inferences in probabilistic networks, especially in Bayesian networks. One system built upon network techniques to implement reasoning at optimum entropy is the expert system shell Spirit [18]. Graph based methods are known to feature a very efficient representation of probability distributions via junction trees and hypergraphs, while MEcoRe works on a model based representation of probabilities. While this is clearly inefficient, the aim of MEcore is to implement subjective probabilistic reasoning, as it could be performed by agents, making various belief operations possible. In particular, it allows changing of beliefs in a very flexible way by taking new, complex information into account. This is not possible with graph based systems for probabilistic inference, as efficient methods of restructuring probabilistic networks still have to be developed.

3 `BT`: Modelling Clinical Brain Tumor Data

For generating an initial knowledge base for clinical brain tumor data we will use various binary and multi-valued variables considering aspects of the patient, the patient's anamnesis, the observed symptoms, the possible diagnosis, etc; a medical justification for these variables and their values along with references to the relevant medical literature is given in [21,22]. Since the prevalence of different tumor types varies with the age of patients, the variable `age` distinguishes patients with respect to the three values `le20` (less or equal 20 years old), `20to80` (between 20 and 80 years), and `ge80` (greater or equal 80 years). The binary variable `warningSymptoms` is true iff warning symptoms like perceptual disturbances or unusual pain in the head are present. Given results of a magnetic resonance tomography (MRT), the variable `malignancy` corresponds to the assumed malignancy of the tumor with respect to the WHO grading system [14]; a higher index corresponds to a higher malignancy. The binary variable `icpSymptoms` indicates whether MRT results provide symptoms for intracranial pressure (ICP). The preoperative physical fitness of patients is evaluated by the ASA (American Society of Anesthesiologists) classification system represented by the variable `ASA`. It is associated with perioperative risks, and a higher value indicates a higher risk. Only the first four states are considered here, as treatment of a brain tumor is of low priority for a higher value. Thus, so far we have:

```
            age : le20, 20to80, ge80
warningSymptoms : true, false
     malignancy : 1, 2, 3, 4, other
    icpSymptoms : true, false
            ASA : 1, 2, 3, 4
```

In `BT`, the ten most common brain tumor types like gliomas and meningiomas [17] are taken into account. Together with the value `other` for any other tumor types, these brain tumor types constitute the values of the variable `diagnosis`:

```
diagnosis : pilocytic-astrocytoma, diffuse-astrocytoma,
            anaplastic-astrocytoma, glioblastoma,
            oligodendroglioma, ependymoma, meningeoma,
            medulloblastoma, cranialnerve-tumor,
            metastatic-tumor, other
```

Finally, there are three variables denoting the therapy, possible complications, and the expected health of the patient. The variable

```
therapy : conservative, surgery, none
```

diagnosis	Adults	Children
glioma		
- glioblastoma	15%	*unspecified*
- pilocytic-astrocytoma	*unspecified*	35%
- diffuse-astrocytoma	10%	*unspecified*
- anaplastic-astrocytoma	10%	*unspecified*
- oligodendroglioma	10%	*unspecified*
- ependymoma	4%	8%
meningeoma	20%	*unspecified*
medulloblastoma	7%	25%
cranialnerve-tumor	7%	*unspecified*
metastatic-tumor	10%	*unspecified*
other	*unspecified*	*unspecified*

Fig. 1. Empirical frequencies of brain tumor types, where *unspecified* stands for rare or unknown (collected from [3,9,15,20])

refers to the therapy to be chosen. We distinguish a conservative therapy without surgery, surgery, or no therapy at all. Possible complications during an inpatient stay are expressed by the variable

$$\texttt{complication}: \ \texttt{1}, \texttt{2}, \texttt{3}$$

which distinguishes the three stages 1 (no complications or minor, completely reversible complications like temporary pain after surgery), 2 (medium or heavy complications with uncertain reversibility like neurological or other functional disorders), and 3 (life-threatening complications like serious internal bleeding or neurological deficits at the risk of brain death). Thus, higher values correspond to more serious complications. The expected health of the patient after inpatient stay is denoted by:

$$\texttt{prognosis}: \ \texttt{very\_good}, \texttt{good}, \texttt{intermediate}, \texttt{poor}, \texttt{very\_poor}$$

The knowledge base BT uses these nine propositional variables as its vocabulary to represent clinical brain tumor data and corresponding expert knowledge. Note that although we have only 9 variables, due to the multiple values they induce $2^2 \times 3^3 \times 4 \times 5^2 \times 11 = 118.800$ possible worlds.

There are various publications containing empirical frequencies of certain brain tumor types. For our initial version of our knowledge base BT, we encode the frequencies given in Fig. 1 that are collected from [3,9,15,20] and that are given relative to the patient being an adult (age=20to80 or age=ge80) or being a child (age=le20). The representation of these frequencies is given by conditionals of the following type

```
(diagnosis=meningeoma|!(age=le20))[0.20]          (3)
(diagnosis=medulloblastoma|!(age=le20))[0.07]     (4)
(diagnosis=cranialnerve-tumor|!(age=le20))[0.07]  (5)
(diagnosis=metastatic-tumor|!(age=le20))[0.10]    (6)
```

where, using the input syntax of MECoRe, ! denotes negation. Additionally, BT contains the probabilistic facts (age=le20)[0.15] and (age=20to80)[0.62] reflecting the age distribution in Germany in the year 2009.

Note that there are some missing frequencies in Fig. 1, and thus, there are no conditionals in `BT` for these missing frequencies. In order to obtain a full probability distribution over all variables and their values, the missing knowledge is completed in an information-theoretically optimal way by employing the ME principle, thus by being as unbiased as possible with respect to each diagnosis with unspecified probability. In MECoRe, the computation of an epistemic state incorporating the knowledge given by `BT` is started by

```
cmd-1:  currState := epstate.initialize(BT);
```

so that `currState` denotes the ME distribution over `BT`.

In order to be able to ask a set of queries instead of just a single query at the same time, MECoRe allows the introduction of an identifier to denote a set of queries. Here, we will illustrate this feature with a singleton set containing an unquantified conditional for the diagnosis under the premise that the patient is older than 80 and that he suffers from warning symptoms

```
cmd-2:  queriesBT := (diagnosis|(age=ge80) ∧ warningSymptoms);
cmd-3:  currState.query(queriesBT);
```

which yields the following probabilities:

diagnosis	*probability*	*diagnosis*	*probability*
`glioblastoma`	0.150	`meningeoma`	0.200
`pilocytic-astrocytoma`	0.035	`medulloblastoma`	0.070
`diffuse-astrocytoma`	0.100	`cranialnerve-tumor`	0.070
`anaplastic-astrocytoma`	0.100	`metastatic-tumor`	0.100
`oligodendroglioma`	0.100	`other`	0.035
`ependymoma`	0.040		

Note that up to now, `BT` does not contain any information about the influence of warning symptoms or the observation that the patient is more than 80 years old. Therefore, in the ME distribution given by `currState`, the corresponding premise given in the queries in `queriesBT` (cf. command line `cmd-2`) does not cause a deviation from the probabilities given in the original conditionals in `BT` and taken from Fig. 1. Note also that the prababilities for the two possible diagnosis values `pilocytic-astrocytoma` and `other` missing for adults in Fig. 1 have also been computed as expected.

4 Revising and Updating `BT`

Besides available statistical data, another important knowledge source is the clinical expert knowledge of a physician. For example, for adults, Fig. 1 tells us that the most frequently appearing glioma tumor type is `glioblastoma`, but no information is provided about its probability given specific symptoms. An experienced physician working with brain tumor patients might state the following conditionals expressing his expert beliefs about the probability of a `glioblastoma` given various observations:

```
(diagnosis=glioblastoma | !(age=le20) ∧ warningSymptoms)[0.20]   (7)
```

```
(diagnosis=glioblastoma | !(age=le20) ∧ icpSymptoms)[0.20]          (8)
(diagnosis=glioblastoma | !(age=le20) ∧ (malignancy=4))[0.40]       (9)
(diagnosis=glioblastoma | !(age=le20) ∧ (malignancy=3))[0.10]      (10)
(diagnosis=glioblastoma | !(age=le20) ∧ (malignancy=2))[0.05]      (11)
(diagnosis=glioblastoma | !(age=le20) ∧ (malignancy=1)[0.01]       (12)
```

Taking into account only Fig. 1, the probability for `glioblastoma` is 15%. Therefore, given the respective preconditions, rules (7) - (9) would increase the probability, whereas rules (10) - (12) would decrease it.

In [21], about 90 conditionals expressing such expert knowledge from a physician's point of view are formulated. With `expertBT` denoting the set of these conditionals, we will incorporate this new knowledge into the current epistemic state. We can achieve this in such a way as if it had been available already in the original knowledge base `BT` by a kind of belief change called *genuine revision* (cf. Sec. 2 and [13,2]). In MECoRe, this is easily expressed by

```
cmd-4:  currState.revise(expertBT);
```

Now, asking the `queriesBT` (cf. command line `cmd-2`) again, the probabilities have changed considerably in the new epistemic state:

diagnosis	*probability*	*diagnosis*	*probability*
`glioblastoma`	0.223	`meningeoma`	0.156
`pilocytic-astrocytoma`	0.050	`medulloblastoma`	0.065
`diffuse-astrocytoma`	0.098	`cranialnerve-tumor`	0.057
`anaplastic-astrocytoma`	0.106	`metastatic-tumor`	0.106
`oligodendroglioma`	0.086	`other`	0.011
`ependymoma`	0.039		

E.g., the probability for `glioblastoma` increased from 15% to 22.3%, while the probability for `meningeoma` decreased from 20% to 15.6%. This is well in accordance with the observations made by physicians working in this area [21].

Now suppose that later on, experts think that the probabilities of conditionals (7) - (9) should be changed to 0.15%, 0.25%, and 0.45%, respectively, and let `gliobNew` denote these three modified conditionals. Genuine revision of the current epistemic state with `gliobNew` would lead to an inconsistency since (7) - (9) and `gliobNew` cannot be satisfied simultaneously. However, MECoRe's update operation of `currState` by `gliobNew` can incorporate the new knowledge in the current epistemic state by choosing the distribution satisfying `gliobNew` and having minimum cross entropy with respect to `currState` (cf. [13,2]). Note that update is the more appropriate operation here, since the shift of the probabilities reflects a changed environment.

5 Prognosis and What-If-Analysis

For the real documented case of a patient being older than 80 years, with `warningSymptoms`, `icpSymptoms`, and `malignancy=4`, asking MECoRe results in a probability of 55.6% for the diagnosis glioblastoma, being very plausible from a physician's point of view. Assuming that glioblastoma were indeed the correct diagnosis and assuming further that a surgery would be chosen, the prognosis for complications that might occur are determined by:

```
cmd-5:  whatIfQ := (complication | (diagnosis | (age=ge80) ∧
                 warningSymptoms ∧ icpSymptoms ∧ malignancy=4));
cmd-6:  hypothesis := ((diagnosis=glioblastoma)[1.0],
                 (therapy=surgery)[1.0]);
cmd-7:  currState.whatif(hypothesis,whatIfQ);
```

Note that what-if is similar to an update except that it does not change the current belief state. The resulting probabilities for complications of grade 1, 2, and 3 are 0.4%, 45.4%, and 54.2%, respectively. While complications of grade 2 or 3 are rare in general, the provided evidence and the given assumptions caused MECoRe to rise the probabilities for these types of complications considerably. After surgical treatment of the given patient, there was indeed a complication of grade 2. From a clinical perspective, the probabilities for `complication` computed by MECoRe is an adequate warning; however, the probability for grade 3 is a bit too pessimistic, since compared to similar patient-risk constellations, life-threatening complications are frequent, but less than 50%. Here, a corresponding adaptation of the conditionals constraining the probabilities for grade 3 complications might lead to a more realistic probability value for this query. Further types of queries for `BT` asking MECoRe for the expected health of patients after inpatient stay, returned a very realistic prognosis from a medical point of view [21]. An example for what-if-analysis where the assumptions are not just facts with probability 1.0 (as in `cmd-7`) is given by `currState.whatif(gliobNew,whatIfQ)`, asking for the probability of `whatIfQ` in the current epistemic state under the assumption that the conditionals in `gliobNew` (cf. end of Sec. 4) hold.

6 Conclusions and Further Work

We reported on a case study using probabilistic logic and the principle of maximum entropy to model clinical brain tumor data and medical expert knowledge in neurosurgery. The knowledge base `BT` contains approximately 110 probabilistic conditionals over 9 multi-valued variables that medical experts identified to be at the core of clinical brain tumor data analysis. Using MECoRe for working with `BT` produced realistic probabilities for diagnosis and prognosis from a clinical physician's point of view. We are currently working on extending `BT`, taking into account additional variables and further refining the medical modelling.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, P.: On the logic of theory change: Partial meet contraction and revision functions. Journal of Symbolic Logic 50(2), 510–530 (1985)

2. Beierle, C., Kern-Isberner, G.: A conceptual agent model based on a uniform approach to various belief operations. In: Mertsching, B., Hund, M., Aziz, Z. (eds.) KI 2009. LNCS (LNAI), vol. 5803, pp. 273–280. Springer, Heidelberg (2009)
3. Bruch, H.P., Trentz, O.: Berchthold Chirurgie, 6. Auflage. Elsevier GmbH (2008)
4. Csiszár, I.: I-divergence geometry of probability distributions and minimization problems. Ann. Prob. 3, 146–158 (1975)
5. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. Artificial Intelligence 89, 1–29 (1997)
6. Dubois, D., Prade, H.: Focusing vs. belief revision: A fundamental distinction when dealing with generic knowledge. In: Nonnengart, A., Kruse, R., Ohlbach, H.J., Gabbay, D.M. (eds.) FAPR 1997 and ECSQARU 1997. LNCS, vol. 1244. Springer, Heidelberg (1997)
7. Finthammer, M., Beierle, C., Berger, B., Kern-Isberner, G.: Probabilistic reasoning at optimum entropy with the MEcore system. In: Lane, H.C., Guesgen, H.W. (eds.) Proc. FLAIRS 2009. AAAI Press, Menlo Park (2009)
8. Gärdenfors, P.: Knowledge in Flux: Modeling the Dynamics of Epistemic States. MIT Press, Cambridge (1988)
9. Hosten, N., Liebig, T.: Computertomografie von Kopf und Wirbelsäule. Georg Thieme Verlag (2007)
10. Katsuno, H., Mendelzon, A.O.: On the difference between updating a knowledge base and revising it. In: Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR 1991, pp. 387–394. Morgan Kaufmann, San Mateo (1991)
11. Kern-Isberner, G.: Characterizing the principle of minimum cross-entropy within a conditional-logical framework. Artificial Intelligence 98, 169–208 (1998)
12. Kern-Isberner, G.: Conditionals in nonmonotonic reasoning and belief revision. LNCS (LNAI), vol. 2087. Springer, Heidelberg (2001)
13. Kern-Isberner, G.: Linking iterated belief change operations to nonmonotonic reasoning. In: Brewka, G., Lang, J. (eds.) Proceedings 11th International Conference on Knowledge Representation and Reasoning, KR 2008, pp. 166–176. AAAI Press, Menlo Park (2008)
14. Louis, D.N., Ohgaki, H., Wiestler, O.D., Cavenee, W.K., Burger, P.C., Jouvet, A., Scheithauer, B.W., Kleihues, P.: The 2007 WHO Classification of Tumours of the Central Nervous System. Acta Neuropathologica 114(2), 97–109 (2007)
15. Müller, M.: Chirurgie für Studium und Praxis, 9. Auflage. Medizinische Verlags- und Informationsdienste (2007)
16. Paris, J.B., Vencovska, A.: In defence of the maximum entropy inference process. International Journal of Approximate Reasoning 17(1), 77–103 (1997)
17. Park, B.J., Kim, H.K., Sade, B., Lee, J.H.: Epidemiology. In: Lee, J.H. (ed.) Meningiomas: Diagnosis, Treatment, and Outcome, p. 11. Springer (2009)
18. Rödder, W., Reucher, E., Kulmann, F.: Features of the expert-system-shell SPIRIT. Logic Journal of the IGPL 14(3), 483–500 (2006)
19. Shore, J.E.: Relative entropy, probabilistic inference and AI. In: Kanal, L.N., Lemmer, J.F. (eds.) Uncertainty in Artificial Intelligence, pp. 211–215. North-Holland, Amsterdam (1986)
20. Steiger, H.-J., Reulen, H.J.: Manual Neurochirurgie. Ecomed Medizin (2006)
21. Varghese, J.: Using probabilistic logic for the analyis and evaluation of clinical patient data in neurosurgery. B.Sc. Thesis, Univ. Hagen (2012) (in German)
22. Varghese, J., Beierle, C., Potyka, N., Kern-Isberner, G.: Using probabilistic logic and the principle of maximum entropy for the analysis of clinical brain tumor data. In: Proc. CBMS 2013. IEEE Press, New York (to appear 2013)