Decision trees optimization for ultrasound detection of fetal abnormalities.

DS3 Summer School 2017

Rémi Besson[1], Erwan Le Pennec[1], Stéphanie Allassonnière[1,2].

CMAP Ecole Polytechnique[1], Université Paris-Descartes[2]

remi.besson@polytechnique.edu

Context and Objectives : Prenatal Diagnosis.

There are three compulsory ultrasound test during pregnancy in France. Some classical measures are done for every women, for example the research of trisomy 21 is well known and mastered. On the contrary there is no strict protocol defined for the research of rare diseases. This is why we want to help obstetricians to improve/systematize ultrasonic diagnostic.

FIGURE 1: An ultrasound view of a 12 weeks embryo.

We should lead the sonographer in order to quickly know which is the disease. It can be viewed as a task of building a decision tree asking for symptoms.

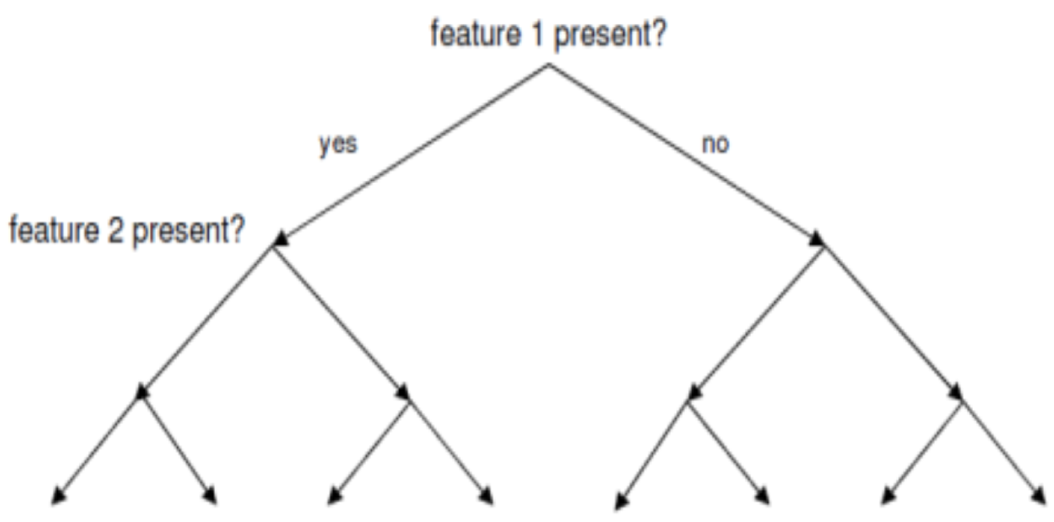

FIGURE 2: A binary decision tree.

Data :

Each disease manifests itself by a combination of abnormalities visible by ultrasound (symptoms). Some symptoms are more probable than others. We know the probability of each disease and the probability of each symptoms knowing the disease. These probabilities come from the literature (acknowledgement to Emmanuel Spaggiari, Physician in Gynecology obstetric at Necker).

Here an extract of our data :

id disease	id symptom	probability of symptom knowing the disease
16	29	0.39
16	136	0.67
16	149	0.50
16	176	0.16
16	181	0.50
16	231	0.75

Notation : $P[A_i \mid S]$ is the probability to have the abnormalities i knowing that we have the disease S.

From Marginals to Joint Distribution via Maximum Entropy Heuristic :

We have $P[A_i \mid S]$ but we need to know $P[A_{i_1}, ..., A_{i_K} \mid S]$.

We face the problem to assign values to probabilities in presence of partial information. We want to add as few information as possible so **we will choose the model with less additional information compatible with our data**.

We use entropy to measure information contained in a distribution, less information meaning more entropy [4].

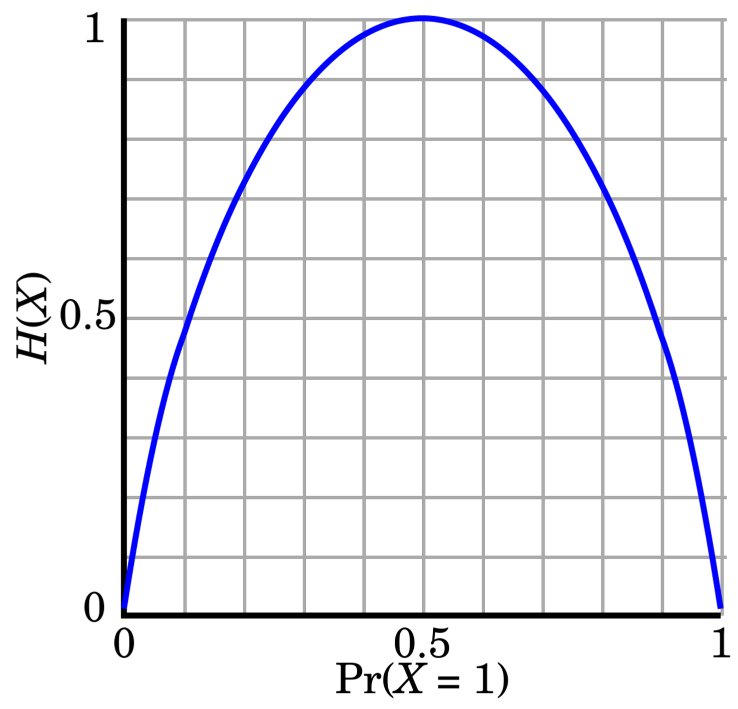

FIGURE 3: Entropy of Bernoulli distribution X.

Then **we choose the model of maximum entropy compatible with our data** (see for example [1]).

We add the information that some cases are impossible (for example the event of having just one symptom isolated is impossible, it is not what we call a disease).

As expected the resulting distribution is something close to the independent one :

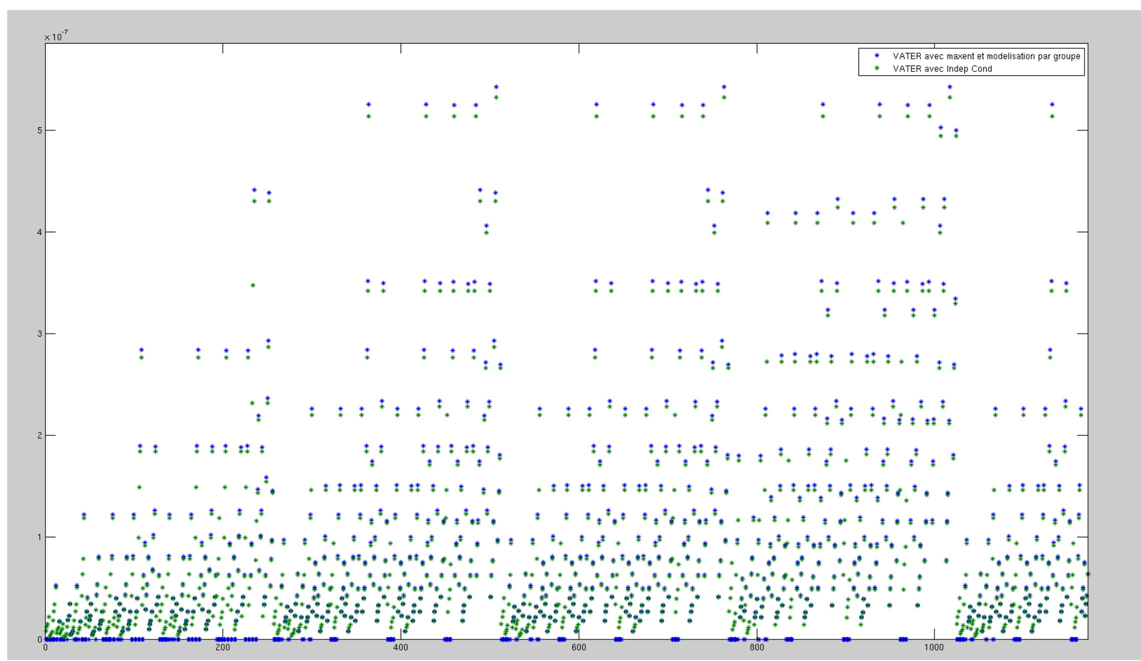

FIGURE 4: Comparison between distribution obtained via maxent heuristic (blue) and distribution obtained via independence hypothesis (green).

If we want to go further we need data. Our algorithm will collect data when it will be used. However our data will never be sufficient to infer the distribution so our model will make a trade-off between entropy and data likelihood. Here is an example with the disease 16, the model without data (maximum entropy with marginal constraints) is in black circle. Now imagine we observed 5 times the cases number 4, 1 times 30 and 8 times 50. The resulting **trade-off between expert data and experience data** is in red stars.

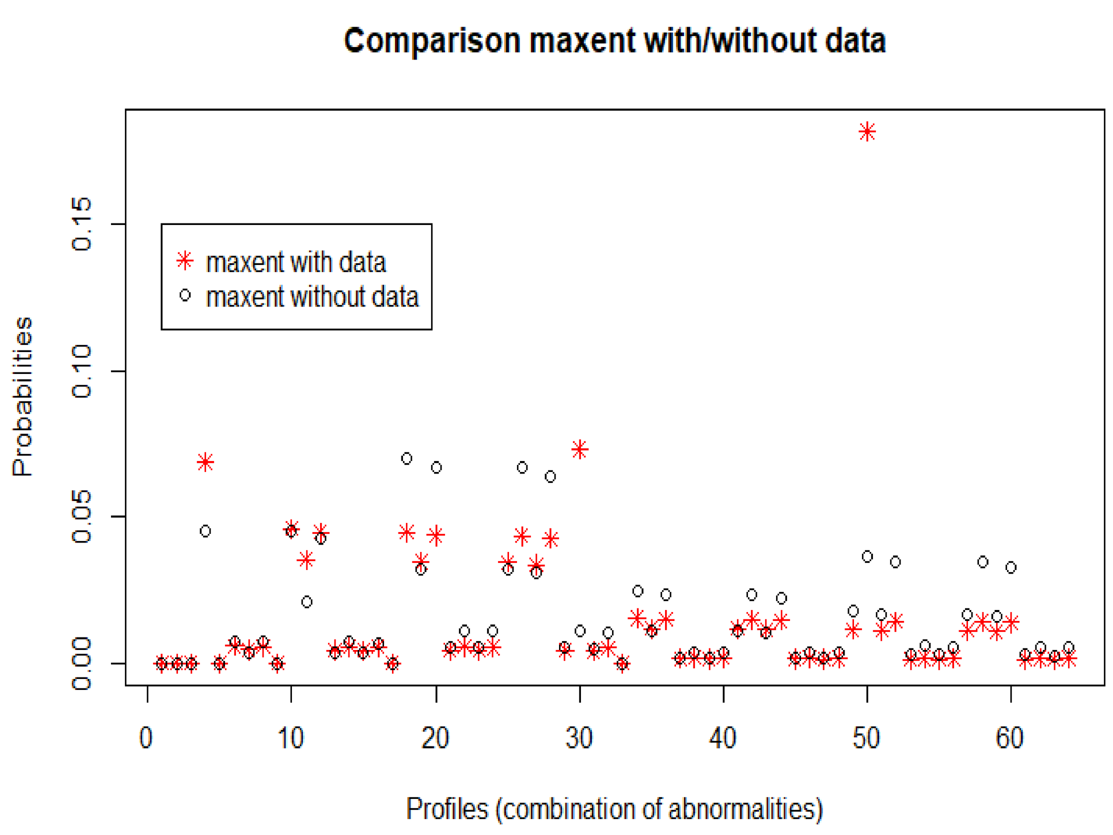

Our algorithm memorise what he observed in his experience considering it more likely to happen again but still consider possible combinations he never observed (just as would do a physician).

Best strategy learning.

Our objective is to determine in which order we have to ask the questions so as to minimize the average number of question necessary to diagnose patient's disease (just as in a 20 question game).

We propose to formulate it in the reinforcement learning framework where someone take actions in a sequential way, receive rewards from the environment and try to maximize his long-term (cumulative) reward [3].

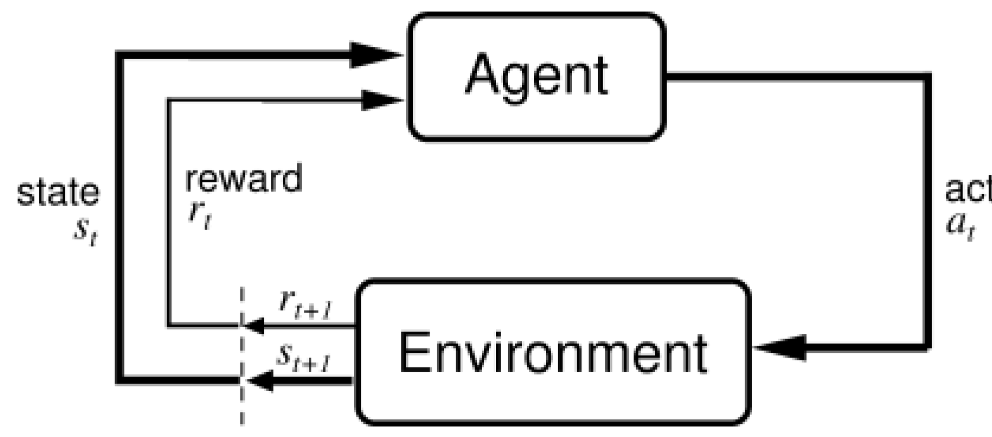

FIGURE 5: Reinforcement learning scheme

Here the agent ask questions about symptoms and each question give him a negative reward : -1. The process stops when it reach a terminal state, a state where we know with high probability which is the disease.

We want to learn a diagnostic policy which associate to each state of knowledge a question to ask.

$$\pi : \mathbb{S} \to \mathbb{A}.$$

We are looking for π^* the best diagnostic strategy, that is the strategy that goes the fastest to terminal states (i.e who maximize rewards).

This kind of problems are well-known and can be solved with algorithm such that Policy Iteration (see [3]). However we face a **problem of high dimension** (300 questions possible) so to use Policy Iteration algorithm or even store π^* is hopeless.

To face this issue we propose a energy-based policy (as in [2]) :

$$\pi(s,a) = \frac{\exp(\theta^T \phi(s,a))}{\sum_b \exp(\theta^T \phi(s,b))}.$$

$\pi(s,a)$ is the probability to take the action a when we are in state s. $\phi(s,a)$ is a vector of features which quantify the interest of taking action a when we are in state s.

The optimal parameters θ^* will be learned doing a stochastic gradient descent solving the problem of maximization of cumulative reward.

Acknowledgment

We are grateful to the team of Profesor Yves Ville (Hôpital Necker Enfants Malades) : Julien Stirnemann and Emmanuel Spaggiari. Thanks to Nicolas Garcelon (Database Engineer), Anne-Sophie Jannot (PhD Inserm), Antoine Neuraz (PhD Inserm) and Anita Burgun (PhD Inserm).

Références

[1] Jaynes E.T. Information theory and statistical mechanics. *The Physical Review*, 1963.

[2] Teh Y.W. Heess N., Silver D. Actor-critic reinforcement learning with energy-based policies. *European Workshop on Reinforcement Learning*, 2012.

[3] Barto A.G. Sutton R.S. *Introduction to reinforcement learning*. Bradford Book, 1998.

[4] Cover T. and Thomas J. *Elements of Information Theory*. Wiley Series, 2006.