

Representing Robot/Environment Interactions Using Probabilities: the "Beam in the Bin" Experiment

Pierre BESSIÈRE, Eric DEDIEU & Emmanuel MAZER

CNRS (Centre National de la Recherche Scientifique)
IMAG/LIFIA (Laboratoire d'Informatique Fondamentale et d'Intelligence Artificielle)
46 ave Felix Viallet, 38031 Grenoble, FRANCE
E-mail: bessiere@imag.fr

Abstract

This paper presents a very simple robotic experiment to illustrate how probabilistic reasoning may be used for sensory-motor systems. We will show how our robot may learn internal representations of its interactions with the environment, how it may predict the sensory result of a given action, how it may generate motor command to reach a wished sensory situation, how it may recognize different status and novel conditions and finally how it may behave consistently to explore and exploit simultaneously its environment.

Keywords

COGNITIVE SCIENCES - AUTONOMOUS ROBOTICS - PROBABILISTIC REASONING

1 Introduction

The goal of this paper is to support, illustrate and discuss the three following thesis:

- 1 - **The activity of a sensory-motor agent is in essence situated.** For both a living organism and a robot the set of all possible values of the signals coming from its sensors and send to its actuators is called the phase space or sensory-motor space of the system. These signals and the corresponding variables may be the results of intricate mechanical, biological, neuronal or computational processes of the raw sensory-motor data. This space is highly structured. When the values of some of the sensory-motor variables have been fixed, the possible values of some other variables are very constrained¹. Any observed dependency between physically independent sensory-motor variables is a relevant information about the interaction between the agent and its environment. Mutually, any pertinent behavior of the agent should cause the observation of such dependencies. As the relevant space is the sensory-motor space, the interaction between the system and its environment is the only mean to learn correct internal

¹ The human articulatory system, for instance, may be described by a phase space made of 11 sensory-motor variables. Seven control parameters (Maeda's model [Maeda88]) are coding the positions of the principal mobile parts of the vocal track. Four variables called "formants" correspond to the maximum energy frequencies of the produced sounds. This space is structured by very strong dependencies between these variables. For instance, when the control parameter "lip aperture" (coding the opening of the lips) is fixed the set of possible sounds (the position in the formantic space) is relatively reduced. At the opposite, given a sound to produce, the possible variation of the lip aperture is very small.

- "representations".
- 2 - **Situated activity is the only mean to deal with environmental complexity.** The central problem of autonomy is to discover, explore, learn, memorize and exploit these dependencies between sensory-motor variables. These representations should be neither symbolic, nor analytic nor geometric. Approximated models of the interaction with the environment are the only way to take into account the complexity of real physical worlds. This kind of internal representations is the only one able to deal with uncertain and incomplete knowledge.¹
 - 3 - **Probability theory is an adequate theoretical foundation and suitable tool to perceive, reason and act in complex environments.** Probability distributions on subspaces of the sensory-motor space are powerful representations of the interactions between the systems and its environment. Many mathematical tools already exist to deal with uncertain and incomplete knowledge and may be used to learn, stock and exploit the dependencies between sensory-motor variables.

To support these three theses the paper will be divided in three main sections.

The first section will discuss the difficulties of the classical approach to robotics. We will show that these difficulties are a particular case of the general "symbol grounding" problem. It will be stressed that the ability to take into account uncertain and incomplete knowledge is a main requirement for autonomous robots.

The second section will introduce the basics of probabilistic reasoning and explain how it deals with uncertain knowledge. It will also state the "principle of maximum entropy" and show how this principle, as a way to take into account ignorance, is a central mechanism to reason despite incomplete knowledge.

The third section will present the "beam in the bin" experiment. This experiment is the most simple one we could imagine with a real robot in a real physical environment. The robot has only one motor variable and one sensor variable. The environment is made of a green plastic dustbin lit by a single lamp. Despite its simplicity this experimental set-up illustrates all the difficulties the classical approach to robotics has to face. Due to its simplicity this experiment demonstrates clearly how probabilistic reasoning may be used in autonomous robotics and why this approach is powerful and interesting.

2 Difficulties of autonomous robotics

The usual approach to robotics may be caricatured by Figure 1.


The designer of the robot has his own human conception of the environment. This conception may be expressed in very abstract terms using analytic, geometric and symbolic descriptions:

- analytic functions to encode the laws of physics;
- geometric language to describe the shapes of the objects and their positions and movements in the 3 dimensions space;
- symbolic notation to name objects and concepts.

The programmer, in some sense, imposes his own model of the environment to the robot when using these terms in his programs.

Everything is fine with this approach as long as the environment is not too complex. For instance, in the field of industrial robotics, the environment of the robot may be controlled and kept

¹ For instance, the analytical model of an ideal mobile robot is very easy to build. Unfortunately, such an ideal robot does not exist. The wheels, for instance, skid and slip on some terrain, for some movements. Consequently, an analytical model of a real robot is impossible to set up. The number of possibly relevant parameters is much too big and we do not know which one are important. For a real robot the value of the different sensory-motor variables is always uncertain and, worst, some of these values are not present.



simple enough to build model using this kind of descriptions.

By "autonomous robotics" we mean precisely, cases where the environment may not be constrained. In these cases the classical technique just described run into some very serious troubles.

The principal difficulty is a very classical artificial intelligence issue known as the "symbol grounding" problem. The robot has to bridge the gap between the raw signals coming from its sensors and the abstract model used in its programs.

In practice, it means that the robot should be able to reason using uncertain knowledge (sensory-motor informations are in essence "noisy") and using incomplete knowledge (clearly, it is not possible to build an exhaustive model of a real physical environment).

We think that answers to these problems may only be found by developing a completely different approach of robotics. The basic idea is to inverse the process by which the programmer imposes his conception of the environment to the robot. We want the robot to build its own "internal representations" of the environment and to propose these "models" to the programmer as the basic programming tools (see Figure 2). In other words, the activity of the robot must be "situated". Performing actions and observing the sensory consequences of these actions the robot should be able to build data structures accounting for the dependencies between the sensory-motor variables.

The fundamental question is then what kind of data structures is adequate to encode these dependencies and is there any satisfactory theory to support the use of this kind of data structures and to propose an answer to the three main identified difficulties:

- treating uncertain information;
- treating incomplete information;
- symbol grounding problem.

The sequel of this paper will mainly try to convince the reader that probability theory should be considered as a possible satisfactory answer.

3 Probability as Logic

Following E.T. JAYNES [JAYNES94] we consider "Probability as Logic" (PaL) as a theory of inference.

By inference we mean simply: deductive reasoning whenever enough information is at hand to permit it; inductive or probabilistic reasoning when - as is almost invariably the case in real problems - all the necessary information is not available. Thus the topic of "Probability as Logic" is the optimal processing of uncertain and incomplete knowledge.

E.T. JAYNES in [JAYNES94]

In this section, given the possible length of this paper we will only try to give a very brief flavor of PaL. The interested reader should read [JAYNES94] for a complete development of this theory and to [BESSIÈRE90] for more details about its implications in cognitive sciences.

3.1 Representation of knowledge

PaL, like classical Logic, works on Aristotelian propositions which could be either true or false.

Logic, starting from a set of true propositions (axioms) tries either to produce deductively new true propositions (theorems) or, given a proposition, to demonstrate that it is true or false. A "state of knowledge" of a logical deductive system is a collection of true propositions.

In contrast, PaL associates a probability of being true or false to all propositions. Instead of deductively producing new theorems, PaL computes the probability of some proposition given the probability of some other ones. A "state of knowledge" of a probabilistic inference system is a probability distribution over the set of propositions.

3.2 Uncertain information

Given some very basic desiderata about plausible reasoning, R.T. COX ([COX46][COX61]) has shown that the only possible way to do inference according to these desiderata is to use probability theory.

Probabilities of propositions are combined with one another according to the two following fundamental rules:

$P(AB X) = P(A X)P(B AX) = P(B X)P(A BX)$	[R1]
$P(A X) + P(\bar{A} X) = 1$	[R2]

Where A and B are some propositions, X stands for the background knowledge, "AB" stands for "A and B", " \bar{A} " stands for "non A" and " $P(A|B)$ " means "probability of A being true knowing that B is true".

It is very easy to demonstrate (see [COX79]) that logic is a particular case of PaL where all propositions have either a probability 0 or 1 of being true. When logic, given "A" and " $A \Rightarrow B$ ", deduces "B", PaL will compute $P(AB|X)$ knowing $P(A|X)$ and $P(B|AX)$ using the rule [R1]. If $P(A|X) = 1$ and $P(B|AX) = 1$ then we obtain $P(AB|X) = 1$ as logic does. However, PaL is obviously more powerful than logic because the previous calculus of $P(AB|X)$ stays perfectly valid even when $P(A|X)$ and $P(B|AX)$ do not take 1 as their values.

In robotics, if we call A some proposition concerning action (for instance A could stand for "The electric intensity on the left wheel motor is 10"), if we call P some proposition about perception and if we call H some hypothesis, PaL may compute:

- 1 - $P(P|AH)$ the probability of observing a given perception knowing that action A has been performed in context H.
- 2 - $P(A|PH)$ the probability of action A knowing either that perception P has been observed or wished to be observed.
- 3 - $P(H|AP)$ the probability of being in the context H knowing that action A has been performed and that perception P has been observed.

3.3 Incomplete information

Given some information, for instance a set of sensory-motor experiences, there are an infinity of probability distributions that may account for these data. However, all these probability distributions are not equivalent. Some appear to be more "coherent", more "probable", more "interesting" than some others. The maximum entropy distribution may be proved to be the "best" of all these possible distributions in some very precise mathematical meaning.

There are several kinds of arguments to justify the use of the maximum entropy principle, intuitive one based on combinatorial considerations (see [JAYNES82]), axiomatic ones (see [SHORE80], [SHORE81] & [SHORE86]) and considerations founded on proofs of coherence between probabilistic inference and conditional probability as expressed by Baye's theorem (see [JAYNES79] & [VANCAMPENHOUT81]). The more precise and convincing arguments are the "entropy concentration theorems" as demonstrated by E.T. JAYNES [JAYNES82] and C. ROBERT [ROBERT90]. They prove that the probability distributions are concentrated around the maximum entropy distributions, moreover they quantify the number of such distributions in a given entropy interval. In fact, they demonstrate that the maximum entropy distribution is the most probable one given the available informations, and the only one that doesn't assume any information that hasn't been given.

Maximum entropy principle is a powerful tool to do inference with incomplete knowledge in the sense that it states what is the "best" state of knowledge (probability distribution) to represent given informations even incomplete.

4 The "Beam in the Bin" experiment

In this section we will present the "Beam in the Bin" experiment. This experiment is based on the most simple sensory-motor device interacting with the most simple real physical environment we could imagine. This simplicity has been a conscious choice. Our goal was not to convince the reader by building a realistic robot. Our aim was, instead, to have an experimental set-up simple enough to present clearly the principles of the proposed approach, to point out the difficulties solved by this method and to ask formally the relevant questions.

4.1 Experimental set-up

The experimental set-up is shown on figure 3.

The "environment" is made of a green plastic bin in which a halogen lamp has been fixed. The




Figure 3: The "Beam in the Bin" experimental set-up

lamp is lightning up the inside of the bin where multiple reflections occur.

The sensory-motor device (robot) is made of a photoelectric cell hanged on a vertical axis. There are two sensory-motor variables:

- θ a motor control variable to specify the angular position of the vertical axis;
- ι a sensory variable corresponding to the light intensity as measured by the photoelectric cell.

We want our device to accomplished the following tasks:

- learn "internal representations" of its' environment;
 - predict value of the sensory variable given the value of the control variable.
 - generate motor command to reach a sensory situation;
 - recognize situations (different environments);
 - recognize novelty (non previously learned environment);
- and finally, behave consistently to exploit and explore changing environments.

4.2 Learning environment "representations" in the sensory-motor space

We firsts place the bin in a given position and fix the intensity of the lamp at a given value. We "explore" this environment using an experimental protocol consisting in drawing at random angles θ , moving the robot to this position and measuring the intensity ι of light on the photoelectric cell. Following this protocol we may build a histogram of the couples of values (θ, ι) observed as shown on figure 4.

The interpretation of this histogram is very easy. We firsts observe high intensity and saturation of the photoelectric cell for angles where the cell faces the lamp. We then observe a secondary intensity peak corresponding to the main reflection for angular positions at 180° of the lamp. It should however be noticed that the histogram does not present a symmetry around the main axis of lightening as one would expect. We will come back on this remark in next section.

4.3 Predicting values of sensory-motor variables

Using "probability as logic" as basic cognitive theory for sensory-motor agent implies that the relevant form for a state of knowledge of a rational agent should be a distribution of probability. Consequently, we want now to derive a probability distribution (a state of knowledge of our robot) from the above histogram (the set of sensory-motor experiences).

To do this we need to give some prior information to our robot. What we want to tell him is that the main factor influencing the intensity ι is the position θ . It means that for a given θ we should always observe the same ι , apart from uncertainties. Consequently, the relevant observables are the mean value and standard deviation of ι for a given θ . Given these observables the principle of maximum entropy tells us that the "best" probability distribution to represent the experimental data is the Normal distribution with the same mean value and same standard deviation than the data.

The state of knowledge (probability distribution) of our robot is the set of Normal distributions shown on figure 5.

Given this state of knowledge, it is trivial to predict the value of the sensory variable ι for a given motor control value θ (D stands for past experiences (ι, θ)):

$$P(\iota|\theta DX) = N(\mu(\theta), \sigma(\theta))$$

$$= \frac{e^{-\frac{(\iota - \mu(\theta))^2}{2\sigma(\theta)^2}}}{\sqrt{2\pi\sigma(\theta)^2}} \quad [1]$$

In the classical approach of robotics, using a symbolic, analytic and geometric explicit model of




Figure 4: The histogram of the observed (θ, ι)




Figure 5: The state of knowledge

the environment the lack of symmetry around the main axis of lightening would be very disturbing. An analytic model of the reflections inside the bin would, for sure, predict such a symmetry. We do not know the exact reason explaining the absence of this symmetry. It could be, for instance, either that the vertical axis of the robot is not exactly the same than the revolution axis of the bin or that the vertical axis of the robot is not perfectly vertical. A robotic engineer would then spend much time and energy to correct this "default". In our approach, given our definition of "autonomy", we do not allow ourselves to constrain the environment and we will not try to reach this ideal fitting between the robot and its environment. In fact, we do not care about this absence of symmetry. We do not need to fit any abstract analytical model. The probabilistic state of knowledge of our robot is sufficient for our purposes and account for the real dependencies between ι and θ even not perfect.

4.4 Generating motor commands to reach a wished sensory situation

The next step is to generate motor commands to reach a wished sensory situation. It is the inverse problem relatively to the previous one.

Solving this problem is the dream of any robotician. It means being able to generate the exact motor command, as Michael Jordan would do, to see the ball going through the basket.

In our obviously much simpler problem, it means being able to generate a motor command θ that gives a good chance to observed a wished i . This is given by:

$$P(\theta | iDX) = P(\theta | DX) \frac{P(i | \theta DX)}{P(i | DX)} \quad [2]$$

The first term $P(\theta | DX)$ corresponds to our experimental protocol to explore the possible angles θ . The exploring policy has been to draw these angles at random and consequently $P(\theta | DX)$ is the uniform distribution:

$$P(\theta | DX) = \frac{1}{N_\theta} \quad [3]$$

The second term $P(i | \theta DX)$ is given by the Gaussian distributions previously learned:

$$P(i | \theta DX) = N(\mu(\theta), \sigma(\theta)) \quad [4]$$




Figure 6: The state of knowledge from above

Figure 7: $P(\theta|(i = 750)DX)$

$P(i|DX)$ may be obtained by normalization, writing:

$$\sum_{\theta} P(\theta|iDX) = 1 \quad [5]$$

We obtain:

$$P(i|DX) = \sum_{\theta} P(\theta|DX)P(i|\theta DX) \quad [6]$$

Finally we get:

$$P(\theta|iDX) = \frac{1}{N_{\theta}} \frac{N(\mu(\theta), \sigma(\theta))}{\sum_{\theta} \frac{1}{N_{\theta}} N(\mu(\theta), \sigma(\theta))} \quad [7]$$

$$= \frac{N(\mu(\theta), \sigma(\theta))}{\sum_{\theta} N(\mu(\theta), \sigma(\theta))}$$

For instance if we want to read an intensity of $i=750$ on the photoelectric cell, we want to compute $P(\theta|(i = 750)DX)$. This corresponds to a cross section of the robot's state of knowledge as shown on figure 6 and will look something like the distribution shown on figure 7.

4.5 Recognizing situations

We now present successively four different environments to our robot. The system learns these environments as described in the previous sections. Its state of knowledge is represented in figure 8.

The first environment is the standard one. The second environment is the same but the bin has been turned of a given angle ϕ . For the third one the intensity of light is lower and, finally, for the fourth environment we removed a small plastic device used to focus the photoelectric cell.

We want the robot to recognize in which one of these four environments it has been placed. To do so the robot goes to a certain θ and measures i . If one such measure is not sufficient, it may do more measures at different θ .

Mathematically we may consider that we have a new variable "e" (for environment) which may




Figure 8: four different environments

take 4 different values. Given a measure (i, θ) we want to compute $P(\varepsilon | i, \theta, D, X)$ given by:

$$\begin{aligned}
 P(\varepsilon | i, \theta, D, X) &= P(\varepsilon | D, X) \frac{P(i, \theta | \varepsilon, D, X)}{P(i, \theta | D, X)} \\
 &= P(\varepsilon | D, X) \frac{P(\theta | \varepsilon, D, X) P(i | \theta, \varepsilon, D, X)}{P(i, \theta | D, X) P(i | D, X)}
 \end{aligned} \tag{8}$$

The robot does not have any prior knowledge about the environment. Consequently, the maximum entropy distribution for the environment knowing the data D is the uniform distribution and:

$$P(\varepsilon | D, X) = \frac{1}{N_\varepsilon} \tag{9}$$

The term $P(\theta | \varepsilon, D, X)$ still corresponds to our experimental protocol to explore the possible angles θ . The exploring policy remains the same and consequently $P(\theta | \varepsilon, D, X)$ is still the uniform distribution:

$$P(\theta | \varepsilon, D, X) = \frac{1}{N_\theta} \tag{10}$$

The term $P(i | \theta, \varepsilon, D, X)$ is still given by the Gaussian distributions previously learned:

$$P(i | \theta, \varepsilon, D, X) = N(\mu(\theta, \varepsilon), \sigma(\theta, \varepsilon)) \tag{11}$$

The denominator is obtained by normalization:

$$P(i, \theta | D, X) P(i | D, X) = \sum_{\varepsilon} P(\varepsilon | D, X) P(\theta | \varepsilon, D, X) P(i | \theta, \varepsilon, D, X) \tag{12}$$

Finally, after some simplifications we obtain:

$$P(\varepsilon | i, \theta, D, X) = \frac{P(i | \theta, \varepsilon, D, X)}{\sum_{\varepsilon} P(i | \theta, \varepsilon, D, X)} \tag{13}$$

Experiments proved that in most cases one measure is enough to discriminate the environments.

However it should be noticed that this process may very easily be applied recursively.

4.6 Recognizing novelty

One drawback of the process described in the previous section is that we have to tell our robot when the environment is changed. This is required because the learning process supposes that the system has collections of data corresponding to a single environment. We would like the robot to be able to recognize novelty of an environment without being told.

To do so, when the robot has learned n environments, we consider that its state of knowledge is made of the n corresponding probability distributions $P(\iota|\theta|(\varepsilon = j)DX)$ ($1 \leq j \leq n$) plus the uniform distribution $P(\iota|\theta|(\varepsilon = 0)DX)$.

$$P(\iota|\theta|(\varepsilon = 0)DX) = \frac{1}{N_\iota N_\theta} \quad [14]$$

With this hypothesis, when some data (ι, θ) corresponding to a "new" (not previously learned) environment, are acquired by the robot the most probable value for ε calculated as described in the previous section turns out to be 0 corresponding to the uniform distribution. This is a very clear and very simple indicator of "novelty".

From [1], [13] and [14] we get:

$$\frac{P((\varepsilon = j)|\iota\theta DX)}{P((\varepsilon = 0)|\iota\theta DX)} = N_\iota N(\mu(\theta, j), \sigma(\theta, j)) \quad [15]$$

which is quite easy to compute.

4.7 Exploiting few experiments

Another drawback of the process first described is that the robot need to achieved completely his sequence of experiments before starting exploiting them. For instance, in the experiment described above the robot had to do $66+100=6600$ measures to build the histogram of a single environment.

We may give him a very useful information by telling him that both the functions $\mu(\theta)$ and $\sigma(\theta)$ seem very regular and smooth. This may be stated as a prior knowledge telling, on one hand, that $P(\iota|\theta DX) = N(\mu(\theta), \sigma(\theta))$, and on the other hand, that when trying to compute $P(\iota|\theta DX)$ for values of θ not yet explored, linear interpolation of the values of μ and σ may be used.

Figure 9 shows the state of knowledge of the robot after only 24 measures when using this prior information. Figure 9 should be compare with Figure 6.




Figure 9: State of knowledge after 24 sensory-motor experiences

4.8 Exploiting and exploring changing environments

In this last experiment, our robot at first a complete ignorant, is placed successively in different environments, and we want it to simultaneously discover, learn and exploit its world. It is given some instructions (wished light intensity ι) by some unknown mechanism. Its goals are to both try to go to a satisfying θ and to learn more about its environment. The used algorithm is the following:

GET-LIGHT(ι)

Begin

- 1 - Compute $P(\theta|\iota, \epsilon_j DX) = \frac{1}{Z} P(\theta|\epsilon_j DX) P(\iota|\theta, \epsilon_j DX)$;
- 2 - Draw an angle θ according to distribution $P(\theta|\iota, \epsilon_j DX)$;
- 3 - Move to θ ;
- 4 - Measure ι ;
- 5 - Compute the probabilities $P(\epsilon_j|\theta, \iota DX)$ for the different environments;
- 6 - Update the parameters of the algorithm;

End

Given a wished ι , in step 1 we compute the distribution $P(\theta|\iota, \epsilon_j DX)$. This distribution is the normalized product of the two distributions $P(\theta|\epsilon_j DX)$ and $P(\iota|\theta, \epsilon_j DX)$. $P(\iota|\theta, \epsilon_j DX)$ are the Gaussian distributions corresponding to the state of knowledge of our robot learn by previous experiences D and using the interpolation rule of section 4.7 for θ not yet explored. $P(\theta|\epsilon_j DX)$ corresponds to our experimental protocol (i.e. our policy to explore environments). For a given θ , the adopted policy is: the more "surprising" the last measure, the more chance to revisit this θ . $P(\theta|\epsilon_j DX)$ is an exponential function of the difference between the last measured ι and the predicted ι as given by $P(\iota|\theta, \epsilon_j DX)$ divided by the variance in θ .

At step 5, we compute the probabilities of the different environment ϵ_j . Three different cases may happen:

- 1 - The current environment is the most probable, the robot keeps on learning and exploiting this environment.
- 2 - Another environment previously learned is the most probable, the robot behaves according to this model and incorporates the new acquired information to this already existing model.
- 3 - The uniform distribution is the most probable one. This is a sign of "novelty" (see section 4.6). If subsequent measures confirm that the robot is in a new environment then the robot creates a new model and starts learning it.

Finally, at step 6 we update the different parameters of the algorithms to take into account the new experience. In particular we update the mean and standard variation of the Gaussian distribution $P(\iota|\theta, \epsilon_j DX)$ corresponding to the angle θ just experienced. Furthermore, we update $P(\theta|\epsilon_j DX)$ given the just measured ι and the predicted one.

This algorithm has five very interesting qualities:

- 1 - There is no need for a learning phase followed by an exploitation phase. The robot simultaneously learns, explores and tries to fulfill its goal.
- 2 - The robot always produces a move. It tries to do its "best" with the knowledge it has. The result is very poor if the environment is not well known but this is inescapable.
- 3 - No supervision of any kind is required.
- 4 - No decision thresholds are used. The moves θ are drawn randomly according to the distributions $P(\theta|\iota, \epsilon_j DX)$.
- 5 - PaL offers a clear theoretical background for this algorithm.

5 Conclusion

This very simple "Beam in the Bin" experiment has been very instructive. We think that it illustrates a very promising and completely new direction to deal with the problems of autonomous robotics.

We are, at present, working in several different directions:

- 1 - We try to recognize and locate different objects put in the bin according to the way their presence's modifies the reflections inside the bin.
- 2 - We aim to make "sensor fusion" using our probabilistic inference framework.
- 3 - Finally, we work on a small mobile robot using a camera and some proximeters as sensors. This robotic device is obviously more realistic than the "beam in the Bin" experiment. One of the main task we are interested in is to learn sensory-motor representations of obstacle avoidance.

References

- [BESSIÈRE90] Pierre Bessière; *Toward a synthetic cognitive paradigm: Probabilistic Inference*; in *COGNITIVA90* edited by T. Kohonen & F. Fogelman-Soulie; Elsevier Science Publishers, 1991.
- [COX46] R. T. Cox; *Probability, Frequency, and Reasonable Expectation*; American Journal of Physic, N°17; 1946.
- [COX61] R. T. Cox; *The algebra of probable inference*; The John Hopkins Press, Baltimore, 1961.
- [COX79] R. T. Cox; *Of inference and inquiry, an essay in inductive logic* in *The maximum entropy formalism*, edited by Raphael D. Levine & Myron Tribus; M.I.T. Press, 1979.
- [JAYNES79] Edwin T. Jaynes; *Where do we stand on maximum entropy?* in *The maximum entropy formalism*; edited by Raphael D. Levine & Myron Tribus; M.I.T. Press, 1979.
- [JAYNES82] E. T. Jaynes; *On the rationale of maximum-entropy methods*; Proceedings of the IEEE, 1982
- [JAYNES94] E.T. Jaynes; *Probability Theory: The Logic of Science*; in preparation; publicly available by FTP at "bayes.wustl.edu".
- [MAEDA 88] S. Maeda; *Improved articulatory model*; JASA (Journal of the Acoustical Society of America); V. 81, N°S1; 1988
- [ROBERT90] Claudine Robert; *An entropy concentration theorem: applications in artificial intelligence and descriptive statistics* in the Journal of Applied Probabilities; september 1990.
- [SHORE80] J. E. Shore & R. W. Johnson; *Axiomatic derivation of the principle of Maximum Entropy and the Principle of Minimum Cross-Entropy*; IEEE Transactions on Information Theory, 1980.
- [SHORE81] J. E. Shore & R. W. Johnson; *Properties of cross-entropy minimization*; IEEE Transactions on Information Theory, 1981.
- [SHORE86] John E. Shore; *Relative entropy, probabilistic inference and A.I. in Uncertainty in Artificial Intelligence*; edited by L. N. Kanal & J. F. Lemmer; Elsevier Science Publishers, North-Holland, 1986.
- [VANCAMPENHOUT81] J. M. Van Campenhout & T.M. Cover; *Maximum entropy and conditional probability*; IEEE Transactions on Information Theory, 1981.