6

Probabilistic Learning

Experience is the teacher of all things.
Commentarii de Bello Civili *(Commentaries on the Civil War), 2. 8,*
Julius Caesar (50s or 40s BC)

There can be no doubt that all our knowledge begins with experience.
Critique of Pure Reason*, B 1, Immanuel Kant (1781; 1787)*

Experience is a brutal teacher, but you learn. My God, do you learn.
Of Other Worlds: Essays and Stories*, Clive S. Lewis (2002)*

6.1 Introduction

An intuitive tell-tale of intelligence is the ability animals possess, particularly humans, of *learning from experience*. So, in fact, when we set out in designing *truly* intelligent systems in robotics, the general aim is to conjure up an architecture that is equally capable of:

- reasoning about the surrounding world given observed data, thereby generating a *representation* – see Chapter 2 to recall what this means in terms of perception;
- learning better representations for the future from the data it is gathering in the present, therefore preparing for *generalisation* – i.e., increasing cognitive performance by refining its internal model of the world as new data becomes available.

A subset of the artificial intelligence research area called *machine learning* is precisely about the construction of such artificial cognitive systems. Machine learning focusses on *prediction*, and the main goal of a learner is to improve its predictive capabilities from new data through generalisation based on what is called the *training data*, generically denoted as Δ. The training data is provided to the system with the hope of reflecting as closely as possible the nature of the model.

Machine learning algorithms are usually classified according either to the input they are provided with or to the desired outcome of their application, of which a partial taxonomy including the most commonly used types of algorithm is provided next:

J.F. Ferreira and J. Dias: *Probabilistic Approaches for Robotic Perception*,
Springer Tracts in Advanced Robotics 91,
DOI: 10.1007/978-3-319-02006-8_6,

Supervised learning – This is the simplest type of machine learning approach. These algorithms are externally guided during a *training phase* so as to generate a function that maps inputs to outputs (also called *labels*, because they are often provided by human experts labelling the training examples). For example, in a classification problem, the artificial learner approximates a function mapping a vector into classes by examining input-output examples of the function.

Unsupervised learning – In this case, the set of inputs is modelled automatically without labelling.

Semi-supervised learning – This method combines both the previous approaches in order to generate an appropriate function or classifier.

Reinforcement learning – In this case, the system learns how to act by observing the impact of each action on the environment, and using this feedback to define a set of rewards.

Probabilistic approaches, and more specifically Bayesian modelling, provide a unifying framework that is able to *explicitly and inherently* address perception, reasoning and learning. As a matter of fact, learning, in this perspective, is "just" another form of probabilistic decision – the cognitive system effectively decides on what its internal model should be considering the training data it is provided with. Let us elaborate further on this notion in the following section.

6.2 Probabilistic Learning as a Decision Process

Learning, from the probabilistic point-of-view, can manifest itself in one of two ways (or generically as a combination of both):

- through the estimation of the *parameters* defining the probability distributions that relate random variables in a model by previously assigning conditional plausibility, when the nature of the relations between variables (otherwise called the *structure* of the model, in an allusion to its graphical representation) is known or defined beforehand;
- through the determination of the structure of the model, when it is unknown.

Assume a generic probabilistic model we wish our robot to learn, represented by $\Pi : \Theta \wedge \Pi'$, where Π' represents the model's *parametric form or structure* and Θ represents the respective *parameters*, providing a domain governed by some underlying distribution P^*. As in section 4.3.3, we are given a set of i.i.d. samples of distribution P^* (i.e., the respective training data), denoted as $\Delta = \{\delta_1, \ldots, \delta_N\}$. We are also given a family of models (i.e., a collection of different structures with arbitrary parameters sharing the same variables in their joint distributions), and the task of our robot is to learn from the training data a model $\tilde{\Pi}$ in this family defining a distribution $P_{\tilde{\Pi}}$