
Fundamentals of Bayesian Inference

Probability theory is nothing but common sense reduced to calculation.

Laplace (1819)

Orthodox thinking, then and now, wants us to define probabilities only as physical frequencies, and deplores any other criterion as not “objective”.

Yet when confronted with a (literally) dirty, objective real problem, common sense overrides orthodox teaching and tells us that to make the most reliable inferences about the special case before us, we ought to take into account all the information that we have, whatever its nature.

“Where do we go from here?”, in Maximum-Entropy and Bayesian Methods in Inverse Problems, p. 21–58, Jaynes (1985)

1.1 Introduction

Since the term *robot* (from the Czech or Polish words *robota*, meaning “labour”, and *robotnik*, meaning “workman”) was introduced in 1923 and the first steps towards real robotic systems were taken by the early-to-mid-1940s, expectations regarding Robotics have shifted from the development of automatic tools to aid or even replace humans in highly repetitive, simple, but physically demanding tasks, to the emergence of autonomous robots and vehicles, and finally to the development of service and social robots.

Along this journey from *automaticity* to *autonomy*, the field of robotics has unavoidably converged towards the field of artificial intelligence (AI) — one has but to notice that the word “autonomous” has its roots in the Greek for *self-willed*. As such, it has suffered the same fate as AI, losing some of its credibility and the power to stimulate the general public’s imagination since the late 1980s. The reason for this is that both AI and robotics set expectations too high as for where autonomy was concerned, aiming for solving difficult cognitive problems through algorithms developed from either philosophical and anthropological conjectures or misconstrued notions of cognitive reasoning, without having yet unveiled even a few of the processes through which biological organisms, in fact, solve these same problems seemingly effortlessly.

One of the major obstacles for reliable robotic autonomy is the problem of the extraction of useful information about the external environment from sensory readings — in other words, *perception*. As of until recently, robots and computer-based devices had been in a state of sensory deprivation. This was in stark contrast with how humans and other natural organisms interact with their everyday environment, efficiently utilising a variety of sensors, such as visual, auditory, haptic, magnetic and odour sensing, just to mention

but a few, even for apparently simple tasks. On the other hand, even when taking into account how these organisms perform using each of these sensory sources individually, their efficiency and efficacy in doing so highly surpasses robotic performance in most cases. Regardless of the unfulfilled expectations, artificial perception has nonetheless witnessed major developments.

Introspection fools us into thinking that perception is deterministic and certain — indeed, as perceptual beings, we rarely question the veracity and accuracy of the understanding of the surrounding world resulting from our senses. As a matter of fact, in psychology the notion of “perception” has often been defined as “a *single unified* awareness derived from sensory processes while a stimulus is present”. This is in accordance with what our common-sense tells us at first glance about what we perceive: “the room that I’m seeing has a chair, a table, and there’s no doubt or uncertainty about that”.

However, recent research on biological perception systems such as the perceptual pathways in our own brain are beginning to question this view. Doya, Ishii, Pouget, and Rao [2] demonstrate throughout their book how an alternative view, the idea that perception is a result of a probabilistic inference process, has surfaced in countless studies in neuroscience, suggesting that the human brain somehow represents and manipulates uncertain information, which can be described in terms of probability distributions.

The implications of these findings, in fact, surface against all of what our common-sense tells us, whenever we are confronted with scenarios which do not conform with what our brains are preprogrammed to accept, as *perceptual illusions* and *bistable percepts*. Nevertheless, the huge amount of natural scenarios which the human perceptual brain is able to cope with (apparently flawlessly) is absolutely astounding. On the contrary, it is particularly striking that robotic perception systems that attempt to tackle more generic and complex problems still cannot rival the performance of a three year-old child due to the lack of adaptive behaviour. Indeed, the amount of restrictions that have to be imposed for robotic perception systems to be able to deal with uncertainty using deterministic and ad hoc approaches is non-negligible (see Fig. 1.1).

When one of the authors of this text was in his teens¹, one day, while reasoning about Prolog programming and trying to replicate the ELIZA program by MIT professor Joseph Weizenbaum, at the end of that day thought “AI is a dead-end: as things stand, the if-then-else rationale and absolute, binary logic just don’t make sense – they sure don’t reflect more than a fraction of *how* we think!”

Aristotle, more or less two and a half thousand years ago, proposed the two famous *strong* syllogisms:

¹ Still dreaming of mobile and humanoid robots in distant planets after reading an issue of “Popular Science” left on his lap by a close friend (nowadays a long-time “partner in crime” in robotics research).