Università Ca'Foscari Venezia

Scuola Dottorale di Ateneo
Graduate School

Dottorato di ricerca
in Informatica
Ciclo XXV
Anno di discussione 2014

Philosophical Aspects in Pattern Recognition Research

SETTORE SCIENTIFICO DISCIPLINARE DI AFFERENZA: INF/01
Tesi di Dottorato di Teresa Scantamburlo, matricola 783915

Coordinatore del Dottorato	Tutore del Dottorando
Prof. Riccardo Focardi	Prof. Marcello Pelillo

Università Ca' Foscari di Venezia

Dipartimento di Informatica
Dottorato di Ricerca in Informatica

Ph.D. Thesis: 783915

Philosophical Aspects in Pattern Recognition Research

Teresa Scantamburlo

Supervisor
Marcello Pelillo

PhD Coordinator
Riccardo Focardi

November, 2013

Author's Web Page: author's home page

Author's e-mail: scantamburlo@dais.unive

Author's address:

Dipartimento di Informatica
Università Ca' Foscari di Venezia
Via Torino, 155
30172 Venezia Mestre – Italia
tel. +39 041 2348411
fax. +39 041 2348419
web: http://www.dsi.unive.it

To my husband Matteo

Abstract

Pattern recognition is the discipline which studies theories and methods to build machines that are able to discover regularities in noisy data. As many of its characterizations suggest, the field is usually considered as a purely engineering discipline. But, in a broad sense, pattern recognition has in fact an interdisciplinary nature. Surprisingly, the attention towards philosophical issues and foundations, which marked the early days of the field, in time has fallen into oblivion. In this thesis, we aim to recover this original interdisciplinary attitude by discussing the philosophical underpinnings of today's pattern recognition research. First we will approach the question of the very nature of the field thanks to the recent developments of both the philosophy of technology and the philosophy of science. This will bring us to consider pattern recognition problems from a slightly different perspective, that is by focusing their relations to some cognitive phenomena (and, in particular, to categorization). Finally, we will undertake a critical analysis of the main research tendencies by making use of Kuhn's notion of a paradigm, one of the cornerstones of the twentieth-century philosophy of science.

Acknowledgments

During my PhD studies I met a lot of great people who enriched my life more than what I could think. So, there would be many I ought to thank but for the sake of brevity I will mention just few of them. As first, I am very grateful to who encouraged me to undertake PhD studies: prof. Carmelo Vigna who supported my training in philosophy and my supervisor, prof. Marcello Pelillo, who taught me to not be afraid of challenging problems and to be curious towards everything. Then I would like to express my sincere gratitude to Viola Schiaffonati who contributed a lot to make this experience an extraordinary resource to my personal, as well as professional, growth. I am also really grateful to Ulrike Hahn who supported me during my period abroad in Cardiff. I really appreciated her active participation and her helpfulness. Moreover I would like to thank Todd Bailey who gave me a lot of precious indications. My gratitude goes also to several mates (in particular to Farshad and Wilayat) and to Samuel Rota Bulò and Nicola Rebagliati for their help and their friendship. Finally I would like to express my deepest gratitude to my husband Matteo who always encourages me and my parents who made this studies possible.

Contents

Preface

This thesis is based on some results achieved during a period of doctoral studies carried out at the Department of Environmental Sciences, Informatics and Statistics (Università Ca' Foscari). As first, it represents the attempt to formulate problems that may often arise in pattern recognition research but that remain nevertheless in the background of the field, or at least of its mainstream. Most of these issues are said philosophical because they are basically concerned with a critical examination of general notions (e.g., that of "science" or "technology") and with a number of further problems widely discussed in philosophy or related areas (such as cognitive sciences). Apparently, the main challenge was to tackle such issues starting from a field, like pattern recognition or machine learning, with a strong attitude towards technicality and applications. Indeed, many efforts were devoted to finding a sensible equilibrium between the philosophical reflection and the real aims of pattern recognition. This resulted in a number of fruitful collaborations, some of which have been discussed in international and national conferences.

In particular, the first chapter stems from the collaboration with Marcello Pelillo and Viola Schiaffonati and includes an examination of scientific and engineering aspects of pattern recognition. The main points discussed in this chapter have been presented at the international conference on History and Philosophy of Computing, held in Paris in 2013 [81]. The third chapter is about the relationship between pattern recognition and cognitive psychology and reports a collaboration with Ulrike Hahn and Todd Bailey during a visiting period in 2012 at the University of Cardiff (UK). Some elements presented in this chapter have been discussed in 2012 at the annual meeting of the Italian association of cognitive science [99]. Finally the fourth chapter, which is concerned with some philosophical assumptions of pattern recognition, is based on a paper written in collaboration with Marcello Pelillo and presented at the conference organized by the Italian association of artificial intelligence in Turin, in 2013. The same paper has been also published as conference proceedings in the Lecture Notes in Artificial Intelligence series [80].

Introduction

"I shall reconsider human knowledge by starting from the fact that we can know more than we can tell."

Michael Polanyi, *The Tacit Dimension* (1966)

In the 1980's Nils Nilsson recommended some important questions to those who are involved in artificial intelligence [72]. He did it stressing the fact that the answer should not be delegated to others outside the field. These questions pertain the nature of artificial intelligence (e.g., "Is AI a coherent subject area? If so, is it a part of computer science, of psychology, or of philosophy; or is it actually an amalgam of parts of these subjects and perhaps others?" [72, p. 1]) and make explicit what in fact contribute to build the inner core of the field. Indeed the development of artificial intelligence, like that of many other disciplines, has been shaped on many different theoretical assumptions. Someone understood AI as the attempt to mechanize the "laws of thought" and ultimately the modern incarnation of what Leibniz had already envisaged in the seventeenth century. A different approach made reference to the Turing test and basically provided an operational definition looking at the behaviour required to perform "intelligent" tasks. Russell and Norvig organized some of these characterizations along two main dimensions, one expressing different degrees of rationality and one definining different point of observations (e.g., reasoning or behaviour), and in so doing ended up with four main definitional styles ("Thinking humanly" - "Acting humanly" - "Thinking rationally" - "Acting rationally") [96]. However, beyond differences, it seems that AI definitions cannot avoid talking about intelligence either as a phenomenon that we can experience or as an abstract problem. Indeed, it is not by chance that a typical AI class starts with some general questions (e.g., what is intelligence? what is learning?) which pertain to different research areas including, among others, epistemology, psychology and more in general cognitive sciences. After all, the development of intelligent systems requires some theories or conjectures on the nature and the operations of the mind suggesting a number of guidelines to the design process.

On the other hand, artificial intelligence may cover a lot of different topics. Over the years, indeed, it has approached a variety of theoretical and technical issues as a cohesive whole interacting with different fields of study. This, for instance, was the experience of the early days of AI when the new engineering challenges were spontaneously correlated to the specific contribution of different scientific disciplines (e.g., biology and physics) and, even more profoundly, to the philosophical investigation. Hence, it is not surprising that many constitutive documents as well as psychological considerations or biological concepts, includes also many philosophical assertions ([66, 114] spring to mind). The same Dartmouth proposal seems to hold that the technological enterprise supposes, in a way, a cognitive demand[1] theorizing that: "every aspect of learning or any feature of intelligence can in principle be so precisely described that a machine can be made to simulate it"[96, p. 17]. Besides the optimism, it was clear enough that speculative knowledge, and specifically understanding the "features of intelligence", could provide the technical endeavour with a good preamble and, more in general, with a type of activity which did not conflict with the engineering skills. In time, this integration remained noticeable just in few areas which, in fact, have become independent from the mainstream of AI research. For instance, the relationship between AI and philosophy has been pursued by the philosophy of mind and in particular by the philosophy of AI, whereas the psychological aspects remained almost exclusively within the domain of cognitive sciences. By contrast, the development of many AI sub-communities (pattern recognition and machine learning are two well known examples) has progressively abandoned the interdisciplinary attitude of early stages making, above all, great technical progress. Cornuejols and Miclet, for instance, observed the same phenomenon analysing some aspects of machine learning curricula: "Overall, there were strong links between Artificial Intelligence and the Cognitive Science fields, where ideas, concepts and empirical studies circulated freely between the different fields. As far as we know, these links tended to slacken during the 90s, following the trends of the nascent "Machine Learning" field towards more mathematics and statistics and more application-oriented focus"[19, p. 2]. Admittedly, the separation between AI research and the study of cognitive phenomena had been already suggested more than a decade ago by J. Hartigan who claimed that, as for classification and clustering problems, "we pay too much attention to the details of algorithms. [...] We must begin to subordinate engineering to philosophy."[41, p. 3].

In this thesis we would like to recover the underlying relationship between science and technology from the inner research development of AI. In particular we will focus on the field of pattern recognition and its relation to philosophy, first of all, shading light on some unexpressed questions and "tacit" ideas underlying the research activity. These elements, indeed, do not arise frequently in scientific and technological disciplines, especially those overspecialised like the fields of pattern recognition or machine learning. However, a lesson we learnt from the philosophy of science of the twentieth century just regards the unavoidable occurrence of theoretical presuppositions in all scientific investigations. The practice of science, indeed, arises from the continuous interaction of different types of knowledge (personal judgements, beliefs, formal concepts, heuristics, etc.). This holds for observations in an experimental set-

[1] Namely it seems to suggest the connection between the "knowing-how" and the "knowing-that".

ting, with respect to which the famous theory-ladenness principle has been stated (see, e.g., Popper, Hanson, Kuhn), but also for the development of a new framework which requires, among other things, the assimilation of what Kuhn calls "metaphysical commitments" [53]. Michael Polanyi went one step further stressing the role of personal knowledge in science whose influence is "passionately and far beyond our comprehension", and acknowledging that "of this responsibility we cannot divest ourselves by setting up objective criteria of verifiability - or falsifiability, or testability, or what you will. For we live in it as in the garment of our own skin"[82, p. 64].

Note that a number of mid-twentieth century philosophers (e.g. Feyerabend, Rorty) took the epistemological pluralism, i.e. the idea that knowledge, science included, is a multifaceted enterprise, in a very radical way so that all criteria of demarcation (attempting to separate science from pseudo-science) fall apart as unjustified and completely useless. Normally, this idea is associated to the theory of epistemological anarchism [33] and to a general rejection of a standard view of rationality, that is the one in which it is possible to draw a line between objective and subjective contents. But epistemological pluralism can be accounted also in a positive way acknowledging the cognitive role of each form of knowledge [62] and, at the same time, tuning specificities by a patient work of cross-fertilization (see, e.g., the effort to develop an interdisciplinary dialogue [62, 32, 71, 68]). Our work takes inspiration primarily by this attitude. That is by the idea that different fields of study can really profit from the exchange of knowledge and from joint efforts to understand complex phenomena.

The present work is organized in the following way. In the chapter 1 we will address the problem of the nature of pattern recognition approaching the question whether this field should be considered science, engineering or both. This will be done by making use of some recent development in the philosophy of science and in the philosophy of technology. In the chapter 2 we will focus on the computational form of pattern recognition problems, that is the typical way in which problems are formulated and solved. The discussion will give also the opportunity to consider the problem of induction, which is a core notion of pattern recognition research, from a philosophical point of view. The connections to the areas of artificial intelligence and machine learning will be also examined. In the chapter 3 we will discuss the comparison between pattern recognition and cognitive psychology, a parallelism which will be carried on thanks to a common term of reference, i.e. the notion of categorization. In the chapter 4 we will address the question whether the field of pattern recognition have achieved the level of maturity in the sense suggested by Thomas Kuhn. This will lead us to analyse the current status of the field with respect to a number of profound commitments. Finally we will propose some concluding remarks.

1
The nature of pattern recognition

Computer science has been plagued since its beginnings by the elusiveness of its very nature, being halfway, as the name itself implies, between science and technology. Dijkstra, for example, insisted on de-emphasizing the role of the machine stressing the intrinsic abstract character of the field; others held that the 'science' in computer science is a misnaming, given its engineering nature [49]. The debate still goes on but, in time, the interdisciplinary nature of computer science has been widely recognized and, accordingly, it is now defined partly as scientific, partly as mathematical, and partly as technological [23]. There are some related fields, however, in which the mutual exclusiveness of the scientific and technological paradigm is still dominant. This is quite evident in several areas of artificial intelligence, such as machine learning and pattern recognition, where only few systematic attempts to understand the interplay between technological and scientific factors have been made. In this chapter, we attempt to approach the question by making use of some recent developments in the philosophy of technology and in the philosophy of science. Our analysis will be focused basically on the research area of pattern recognition and it will be complemented by some concrete examples.

Our discussion will advocate that pattern recognition is a suitable example of the symbiotic relationship between scientific and technological efforts. If we look at the history of the field indeed we observe that most technological progress springs from very scientific issues and early attempts tried not only to provide feasible solutions, but also to uncover the structure of the problems. The case of neural networks is paradigmatic, as their formulation was been clearly inspired by scientific purposes, that is, by the wish of studying and imitating the brain but, in the phase of their renaissance, technical matters prevailed. Indeed, with the (re)invention of the back-propagation algorithm for multi-layer neural networks and, above all, thanks to the impressive results obtained by these new models on practical problems such as zip code recognition and speech synthesis a new wave of excitement spread across the artificial intelligence community. At that point, however, it was widely accepted that these models had no pretension of being biologically plausible except of being interesting computational devices [77]. Bayesianism is another interesting example of the gate allowing pattern recognition to move from theoretical issues to more practical aims. Introduced as a theory, which can characterize the strength of an agent's belief, it provided many inference algorithms with a practical machinery. On the other hand, recent advances in density estimation techniques, such as nonparametric Bayesian methods, have been successfully applied to approach a variety of cognitive processes [97]. This choice is typically useful in problems suffering from a combinatoric explosion and particularly suitable to bridge the gap between the computational and the algorithmic levels of rational models of cognition.

1.1 Pattern recognition between science and technology

Pattern recognition is very hard to define. Indeed, many intellectual activities can be grouped under this term including, among others, categorization and decision-making. Furthermore, pattern recognition has been properly applied to many other natural phenomena (consider, for instance, the processes of pattern formation studied in biology or chemistry [4]) and in the last century it has been extended also to the world of machines thanks to the great development in computer science and other related areas such as cybernetics and artificial intelligence.

In artificial intelligence a popular definition of the term is presented by Bishop in his influential work "Pattern recognition and machine learning". Bishop claims that "the field of pattern recognition is concerned with the automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions such as classifying the data into different categories" [12, p. 1]. But if we extend a bit our examination we will find that there exist several ways to approach the subject and at least two of them are rather different.

On the one hand we could find a "narrow" account which considers essentially a particular interpretation of pattern recognition leaving aside other research perspectives. Usually, the aspects emphasized are mostly engineering in nature and are put forth by reason of convenience rather than of principles. Note, however, that the narrow viewpoint represents the mainstream of today's pattern recognition research even though it is not made always explicit.

On the other hand, we could notice a "broad" perspective which tends, on the contrary, to consider more aspects than those required by a technical challenge. In this respect, the "broad" standpoint has a profound interdisciplinary inclination because it generally brings research on discipline boundaries intersecting knowledge from different fields of study. Unlike the narrow view, it is generally less common in the community research although it has contributed a lot to the development of the field.

Now, let us consider some examples drawn from the main literature so as to make our rubrics more precise. Note, however, that we are not going to provide any strict classification since our main purpose is in fact to point out two general ways of conceiving the discipline.

Normally the "narrow" approach defines pattern recognition as a engineering problem, that is in terms of methods and solutions. Examples of this type abound in the literature and some of them are listed below.

- "It is now abundantly clear that pattern recognition is an immensely broad subject, with applications in fields as diverse as handwriting and gesture recognition, lip-reading, geological analysis, document searching, [...]; it is central to a host of human-machine interface problems, such as pen-based computing." [25, p. xvii]

- "Thus, pattern recognition, or decision-making in a broad sense, may be considered as a problem of estimating density functions in a high-dimensional space and dividing the space into the regions of categories or classes." [37, p. 2]

- "Pattern recognition (PR) concerns the description or classification (recognition) of measurements. PR capability is often a prerequisite for intelligent behaviour. PR is not one technique, but rather a broad body of often loosely related knowledge and techniques. PR may be characterized as an information reduction, information mapping, or information labelling process." [100]

As these characterizations may suggest, the narrow perspective is more focused on the procedure which could successfully perform a pattern recognition task (for instance "dividing the space into regions"), on the form of a feasible solution and on the ways in which all these elements could be measured and compared. Hence, the underlying motivation is always related to concrete context of application like "geological analysis" or "document searching". This fundamental attitude leads researchers and practitioners to overlook the general issues which underlie a pattern recognition problem and which may occur in the study of several phenomena (cognition, evolution, social behaviour, etc). Nevertheless the narrow formulation does not really ignore the great complexity of pattern recognition problems and the profound relationship between the technological effort and other types of investigation (e.g., in the fields of psychology or philosophy, etc). But such a relationship is at most alluded and so often most of knowledge and techniques appear "loosely related."

As for the "broad" point of view, we find out that pattern recognition is depicted more extensively as a general problem that, before dealing with the design of effective techniques, has to do with the study of natural processes. Two examples of this style are the following.

- "Pattern recognition can be characterized as discerning one thing or one form in a field of perception consisting of many parts or examples. [...] In all processes of pattern recognition, the mental activity of unification, or integration, as modern psychologists call it, is at work." [120, p. 7-8]

- "Thus pattern recognition is the decision-making process that assigns to some experiences (carved by this very decision out of the total flow of experience) some internal meanings."[115, p. 365]

Usually, the "broad" picture of pattern recognition points out the general terms of the question emphasizing the role of some abstract dynamics (e.g. "integration" or "seeing-the-one-in-many" [1]) and making reference to a broad spectrum of mental activities (e.g. assigning meaning to experience). Note that encompassing several aspects this approach takes a strong interdisciplinary imprint since much of its contribution relies indeed on different fields of study. A great exemplar of the broad perspective is easily found in the seminal work by Watanbe who regarded pattern recognition under many different guises showing consideration for perception and categorization as well as statistical decision-making and structure analysis.

So considered, the field of pattern recognition intersects a number of different issues, which may go from engineering to philosophy, and combines the design of data analysis techniques with more traditional scientific activities, such as building models, developing theories or simply analysing conceptual notions. Sometimes this resulted in the development of computational models of cognitive abilities [115, 64, 7, 89] and at other times in a supplementary reflection upon the pure technological endeavour [50, 105, 94, 29, 77, 31, 28]. In particular, during the last few decades we noticed an increase of interest towards the scientific attitude of pattern recognition [105, 31, 28, 81] generally expressed by its "natural" inclination to model (and therefore to work as) an empirical phenomenon and specifically the ability to abstract or generalize from observations.

The broad and the narrow perspectives indirectly gave rise to two different characterizations of the disciplinary status of pattern recognition, that is the scientific and the engineering views, respectively. Note that the development of these points of view did not solicit an extensive discussion like the one occurred in computer science or in artificial intelligence, so that the general representation of pattern recognition as a specific field of study (i.e., its disciplinary status) appears rather controversial. Specifically, theoretical and practical issues are often perceived in contrast to each other and, thus, the engineering and the scientific perspective are hardly presented as two mutually related, as well as equally important, aspects of pattern recognition.

On the question whether it makes sense for pattern recognition to adopt a scientific point of view, we cannot avoid to mention the clear position of one of the fathers of the field, Theo Pavlidis who advocated a purely engineering perspective. Indeed, according to Pavlidis "pattern recognition is engineering because we try to design machines that read documents, count blood cells, inspect parts, etc. We must understand the "physics" of the problem and select from amongst available tools the ones appropriate for the problem. It is futile to look for general mathematical/computational techniques that can solve all problems" [77, p. 5]. A similar

[1] This is another beautiful expression used by Watanbe in [120]

position has been recommended specifically for unsupervised classification in a recent paper by von Luxburg et al. precisely to suggest that "clustering should not be treated as an application-independent mathematical problem, but should always be studied in the context of its end-use" [118]. On the other hand this idea seems to be contrasted by the attempts comparing pattern recognition with physical science [105] and more in general by the belief that a scientific approach is not only plausible but even beneficial [31], not to mention the famous axiomatization proposed by Kelinberg [51]. Note that with respect to the field of machine learning the identification with science is even more apparent when we consider the influential work by Pat Langley [57, 58] and the recent studies on scientific discovery [113, 3, 40, 35]. And yet, if we add the fact that the meanings of "science" and "technology" or "engineering" remain implicit for the most part [2] we will probably come to the conclusion that the general picture of the field is a bit confusing and in need of some elucidation.

The scenario presented above suggests us that some instances which are implicitly supposed by the narrow and the broad points of view need more considerations. In particular we think that the engineering and the scientific characterizations of pattern recognition should be understood in depth because the current way to approach the field seems to be strongly unbalanced. Indeed, once the broad approach lost its popularity [3], the research community put forth essentially a narrow account of the field and the opportunity to develop different facets of pattern recognition became a rare case [31]. Nowadays the situation is rather marked by a sort of opposition as the majority of the descriptions get stuck on one of the two components, i.e. the engineering or the scientific, without regard to the strong relationship between these two. In the next sections we will suggest that the contrast between science and technology stems from an oversimplified view of their mutual relationship and that the field of pattern recognition, differently from what we could be used to think, is indeed a suitable example of this cooperative interaction.

1.2 Science *versus* technology

The typical way the dichotomy between science and technology has been conveyed over the years looks like a "superior-subordinate relationship". Hans Poser described this unfair condition without mincing words for sure. He claimed, for instance, that philosophers in most cases "prefer to discuss the rationality of *animal rationale* instead of the products of *homo faber*" whereas scientists generally "look down

[2] In the sense that they are not used to explicitly discuss what science or engineering is expected to deal with in general and therefore in which sense pattern recognition can be considered a scientific and/or engineering.

[3] Indeed works inspired by a broad point view, like the Watanabe's one, are also related to a different way to approach research. Today's criteria of evaluation (based essentially on the number of citations) would probably create not few difficulties to researchers who aimed at more interdisciplinary projects.

on technology as a kind of science-less application of science" stating further that "only if they need some sophisticated new measuring instruments do they accept technology as an auxiliary science" [86]. Actually in the history of philosophy the technical activity has not always been seen negatively [4] although the preference for pure thinking persisted over the centuries. In the main we do not know whether technology was indeed the "Cinderella" of knowledge but we can certainly concede that it has been understood as a merely application of science for a long time. For a general overview on the "spectre" of technology as "a subordinate exercise" we refer the interested reader to [34, 60].

Indirectly the philosophical reflection contributed to such a result. Indeed the early speculations of the philosophy of technology were much more concerned with a number of issues living outside rather than inside technology. The analysis was focused essentially on the social impacts of technology on the society with less regard to technology itself. These studies were undeniably beneficial to understand the complex relations between technology and society but, on the other hand, the intrinsic nature of technology was for the most part ignored (and therefore viewed more as a "black box"). A different approach has been developed since the 1960s and nowadays it is well known under the label "analytical philosophy of technology." Unlike the the pre-existing studies which were in fact closer to social sciences, the analytical approach "regards technology as a practice, basically the practice of engineering. It analyses this practice, its goals, its concepts and its methods, and it relates these issues to various themes from philosophy" [34]. Note that even if throughout the work we will use systematically the term "philosophy of technology" we in fact refer exclusively to the analytical approach.

1.3 The interplay between science and technology

Saying that between science and technology there exists a strong interplay does not mean to remove the differences at all. Indeed even those who questioned the subordination of technology with respect to science generally have never aimed at creating one uniform camp in which both science and technology lose their identity. On the contrary, science and technology are usually presented as grounded on two distinct questions. The former would spring from the study of *what is* whereas the latter would arise from the concern for *what is to be* [108]. In his renowned book "The sciences of artificial", Herbert Simon recalled such a distinction in a slightly different way by noticing that science is concerned with *how things are* while engineering is

[4] For instance Aristotle thought that technique was more than imitating nature. As he put it "generally art in some cases completes what nature cannot bring to a finish, and in others imitates nature" (Physics II.8, 199a15). And even Plato, who maintained that technique consists in imitating nature, applied interestingly a technological image in describing the world as the work of an artisan (see the role of Demiurge in the *Timaeus*). Then much more evident is the appreciation during Renaissance as the Bacon's reflection attested (see for instance the utopian vision of the world described in *New Atlantis*).

concerned with *how things might be* [42]. More expressly the differentiation rests on the underlying intentions. On the one hand science aims at knowing the world (its structure, its regularities, etc.), on the other hand technology aims at producing concrete results (procedures, artefacts, etc.). Therefore it is often said that science is characterized by a noticeable cognitive attitude (the so called knowing-that or knowing why) while technology is more pragmatic (or more interested in the "knowing-how"). Note that this idea in principle emphasizes two characterizing traits: the basic contemplative demand of science, typically expressed by a disinterested conception of knowledge (or what Greek philosophers called *episteme*), and the engineering concern for applications with its methodical search of efficiency. In a way these qualities could be thought as the two focal points of an ellipse. They are representative of two specific identities but at the same time they are joined by a sequence of infinite adjacent points. Therefore, once the foci of our ellipse have been outlined we would like to explore in more details the nature of these connecting points. In which terms we could say that science and technology are continuous to each other?

The relationship between science and technology could be shaped in several ways. The most basic one is a direct consequence of the ancient Greek notion of *techne* which is often imprecisely translated as "art". The term *techne*, however, bears more resemblance to *episteme* than to aesthetic. Indeed what really underlies the notion of *techne* is the same search for the why and hence the same contemplative demand which characterizes *episteme* as well. As Agazzi put it: " The Greek idea *techne* expresses a demand for a theoretical awareness which, so to speak, justifies conceptually that practical knowledge which is already established empirically. *Techne* consolidates this practical knowledge and affords it a certain extension - due to the inherent generality of theoretical knowledge - but is not bound to produce new know-how, or to improve its efficiency" [5, p. 4]. To this end the suffix "ology" that we find in the term "technology" makes this explicit by suggesting a clear theoretical component which provides the reasons of a certain knowing-how.

Until the seventeenth century, however, the general conception of technical activities was undermined by the strong Aristotelian distinction between the celestial region (i.e., "supralunar region"), which is the realm of certain and incorruptible knowledge, and the terrestrial region (i.e., "sublunar region"), which is, by contrast, the world of contingencies of everyday life whose knowledge is always imperfect and subject to change. Therefore, even though there could be a positive relationship between the notions of "episteme" and "techne" (since technical skills might involve the knowledge of causes in the same way as theoretical activity does), the practical work remained "outside" of what was considered as the hard core of theoretical knowledge. Things changed radically with the rise of modern science, that is when knowledge acquisition became inextricably tied to experimental practice and the development of technical tools.

Actually, the practical dimension of knowing started being explored since the time of Renaissance in a clear connection to the idea of human primacy over the nature, typical of that age. Within such a perspective human knowledge started to be exercised through the use and manipulation of nature. In this way intellectual activity "it became strongly allied to the idea of a useful knowledge which would help humans to dominate Nature and to establish a supremacy which would guide and advance practice rather than merely reflect upon nature" [5, p. 5]. Hence it is not surprising that various authors of the Fifteenth century showed great appreciation for technical skills. For example, the erudite philosopher Juan Luis Vives (1492-1540) acknowledged that science should not be controlled only by philosophers or logicians but even by engineers who may know it better than those [95, p. 37]. Then, further appraisal was suggested by essays on mechanics or arts and translations in vernacular of many Latin works about architecture (e.g., Vetruvio's writings) which really abounded throughout the Fifteenth and the Sixteenth centuries.

Later, Scientific Revolution, as well as overthrowing the classical ideas concerning the nature of the universe and the explanations of what occurs within it, introduced new important methodological innovations on which the marriage between science and technology has been ultimately grounded. Starting from this period science became experimental as much of scientific investigation were conduced through experiments, i.e. by means of a *controlled experience* in which theoretical knowledge, the employment of specific tools (e.g., measuring devices) and practical skills combined fruitfully and led to many discoveries. The new science in turn provided a harvest of detailed knowledge that could be used to express the human genius and to develop a new form of knowledge, i.e. a knowledge that was sought for the sake of some technical application.

In time the scientific and technological practices became more and more clearly intertwined. Indeed it has been widely acknowledged that experimental science cannot work without the support of technology and on the other hand the theoretical contribution of technology could be on a par with theoretical research of "ordinary" science. Such analogies were stressed specifically in the philosophy of experimentation [5]. Hans Radder, for instance, stated that "Experiments make essential use of (often specifically designed) technological devices, and, conversely, experimental research often contributes to technological innovations. Moreover, there are substantial conceptual similarities between the realization of experimental and that of technological processes, most significantly the implied possibility and necessity of the manipulation and control of nature. Taken together, these facts justify the claim that the science-technology relationship ought to be a central topic for the study of scientific experimentation" [91, p. 4].

This interplay has been well captured by the recent developments in the philosophy of technology and in the philosophy of science and some of these contributions are provided in the next sections. Specifically we will present some conceptual tools

[5]Even the educational training for aspiring scientists and engineers is indicative of today's integration between science and technology. In the majority of cases the curricula are, indeed, largely identical in the early stages while diverge gradually in rest part.

which can suggest, on the one hand, the cognitive attitude of technology and, on the other hand, the practical dimension of science.

1.4 The contribution of the philosophy of technology

The first great contribution to the study of technology comes from historians. They were among the first to approach technology closely carrying out detailed studies on the genesis and the structure of technological innovations. However, during the second half of the last century the debate on the nature of technological activities solicited a variety of philosophical contributions and specifically the development of conceptual tools which remarkably improved the general understanding of technology as a form of knowledge. For instance in the 1960s three special issues of *Technology and Culture* have been devoted to rethink the relationship between science and technology and in particular the traditional idea of technology as applied science[6]. But even later much has been done to deepen the special characters of the practice of technology emphasizing, for example, the central role of the design process as a structured series of steps which includes important creative elements as well as practical constraints [116, 110]. Note that in the design process what really counts is practical rationality, namely, the ability to provide the criteria on how to act, given specific circumstances. Generating and selecting a list of actions are indeed perceived as a crucial aspect of a technological activity and ultimately a reason of its cognitive side. Bunge somehow encapsulated such ideas in his notion of "operative theory" which, we think, could well characterize most of the practical rationality incorporated by pattern recognition research.

1.4.1 Bunge's operative theories

Bunge was one of the authors who suggested a sense in which technology may be considered as applied science, but without blundering into the conventional idea of subordination. According to him, indeed, technological theories might be divided in two main classes, the substantive and the operative theories. The main difference between them is that "substantive technological theories are always preceded by scientific theories, whereas operative theories are born in applied research" [16, p. 331]. Therefore in technology there are some theories which are genuinely an application of some pre-existing scientific knowledge, but there are as well many others which are in fact originated by the technological context itself. Operative theories express the most creative side of technology as opposed to the one which conveys more the passive role of technology. Hence, it is easy to understand the relevance of such theories in order to shed light on the theoretical contribution of technology.

[6] Some works are [16, 108, 52].

The distinctive point of operative theories is their direct connection to the context of action. Indeed, as Bugne put it "a theory may have a bearing on action either because it provides knowledge regarding the objects of action, for example machines, or because it is concerned with action itself, for example, with the decisions that precede and steer the manufacture or use of machines" [16, p. 331]. Unlike substantive theories, the operative ones are strictly dependent from the practical problem they are born to. So, while the application of the relativistic theory of gravitation to the design of a spaceship launch is a sort of technological substantive theory all theories of decision, planning and optimizations are suitable examples of operative theories. Indeed queuing models or theories of airways management make little if any use of substantive knowledge supplied by other sciences such as physics or biology.

In this respect operative theories do not differ so much from the theories of science. Many characteristic traits of theorizing in science are indeed presented even by operative theories. For instance, like scientific theories they generally make reference to some idealized models of reality and employ a great deal of theoretical concepts (e.g., probability, causation, generalization, etc.). Yet, if there is a proper training (e.g., the system is provided with empirical information) they could produce predictions or retrodictions of events and accordingly undergo various empirical tests. In so doing, Bunge ultimately suggests us to consider technology on par with pure science and operative theories as the exercise of scientific method in the solution of practical problem. But at the same time he maintains that the epistemological relation between science and technology is asymmetrical since pure theory can ground successful practice but practical successes cannot ground pure theory. Indeed Bunge specifies that "the practical success or failure of a scientific theory is no objective index of its truth value" [16, p. 334] as what technology primarily aims to is the achievement of practical goals and this in general does not require a fully consistent theory. Interestingly the same idea seems to be mirrored by the Laudan's account of empirical problems and thereby required by those investigations which look at how science and technology overlap.

1.5 The contribution of the philosophy of science

Formally, the philosophy of science claims a much longer tradition compared to the philosophy of technology. But in fact its contribution to an integrated view of science and technology is rather new and this delay may be reasonable ascribed to the early indifference towards technology itself we have already discussed. Yet, what really increased the growth of such a novel perspective is the influence of the history of science in many philosophical studies in particular during the 1960s and 1970s. And probably Kuhn's contribution is one of the most famous example of the prominence of history to the analysis of scientific changes. Kuhn, indeed, in the very beginning of his well-known book says clearly that "history, if viewed as a repository for more than anecdote or chronology, could produce a decisive transformation in the image of science by which we are now possessed" [53, p. 1]. In time Kuhn's analysis has been fruitful even to explore the analogies between technology and science and this can be easily verified by looking at the wide application of Kuhn's model within technology but, e.g., also within economics and social sciences [7]. As well as Kuhn many other philosophers attempt to model scientific changes through the lens of history and with profitable results even to the study of technology. Some interesting examples are given by Lakatos and Laudan. In particular the latter explained science as the result of a problem-solving activity whose progress is legitimate by the achievement of specific ("non-transcendent") goals. Even though Laudan's proposal could be critical if applied to the entire spectrum of scientific activity [67] we think that it may be suitable, on the other hand, to understand how science moves towards technology and, in our specific case, how science may be carried out in pattern recognition.

1.5.1 Laudan's problem solving approach

According to Laudan "science is essentially a problem-solving activity" [59, p. 11] and its characterization can perfectly avoid any reference to the tiresome questions of truth and ontology. Even though many earlier philosophers of science exploit the idea of problem-solving (e.g., Peirce, Dewey, Kuhn), Laudan thinks that the very nature of (scientific) problems and that of their solutions have been largely ignored. Specifically he aims at providing more details about the very ramification of this approach by scrutinizing the types of problems and the possible relationships between problems and theories.

One important distinction in Laudan analysis is between empirical and conceptual problems. Empirical problems are "anything about the natural world which strikes us as odd, or otherwise in need of explanation" [59, p. 15]. Popular examples of empirical problems are those raised by events that present impressive regularities (e.g., the fall of heavy bodies towards earth or the effect of certain combined elements, etc.). While conceptual problems are "higher-order questions about the well-foundedness of the conceptual structures (e.g., theories) which have been devised to answer the first-order questions" [59, p. 48]. Usually they can spring from internal inconsistencies, that is, when a theory exhibits incoherences or ambiguities, or from external conflicts, for example, when a theory "makes assumptions about the world that run counter to other theories or to prevailing metaphysical assumptions" [59, p. 146].

What we would like to consider now is the essential character of the empirical problems. They are called empirical because we treat them as if they were problems about the world. This means that an empirical problem should not be equated with the notion of a fact. Laudan argues that what matters for being an empirical

[7]For a general overview of the application of Kuhn's model to different fields of study see [39]. For a specific application to technology see [18, 124] and to artificial intelligence see for instance [102, 43, 22]. Note moreover that even the last chapter of this thesis will exploit a kuninan analysis.

problem is to be believed to be a fact. In the history of science there were indeed several examples of facts which did not pose any problems and problems which turned out to be in fact counterfactual [8]. Note that within this perspective what is shaped by engineers or computer scientists in a design process can be properly considered an empirical problem as it is thought to be an actual state of affairs.

Another basic consideration regards the rationale of problem solution. According to Laudan the logic of "solving a problem" is rather different from the logic of "explaining a fact". Laudan, indeed, associates explanations to a requirement of exactness which would be lacking in the pragmatics of solutions. Accordingly a theory solves a problem "so long as it entails an approximate statement of the problem" [59, p. 22]. This approximate character of problem solution is found in several experimental experiences ("Newton was not able to explain exactly the motion of the planets, Einstein's theory did not exactly entails Eddington's telescopic observations" [59, p. 23]). Moreover the question whether a theory solves or not a problem is fundamentally independent from the truth of the theory itself ("Ptolemy's theory of epicycles solved the problem of retrograde motion of the planets, regardless of whether we accept the truth of epicyclic astronomy"[59, p. 24]). This suggests ultimately that in appraising the merits of theories, it is more important to ask whether they are adequate solutions to the focused problem than it is to ask whether they are "true" or "well-confirmed" [9]. Laudan, indeed, noticed that the history of science abounds with theories which were considered perfectly adequate at a time and hopelessly inadequate at another one and that much of these changes depended on the evolution of criteria for the assessment of solution acceptability.

1.6 Science and technology in pattern recognition

If we look at the history of the field, we observe that most of the technological progress springs from very scientific issues and early attempts tried not only to provide feasible solutions, but also to uncover the structure of the problems. The case of neural networks is paradigmatic, as their formulation was been clearly inspired by scientific purposes, that is, by the wish of studying and imitating the brain but, in the phase of their renaissance, technical matters prevailed. Indeed, with the (re)invention of the back-propagation algorithm for multi-layer neural networks and, above all, thanks to the impressive results obtained by these new models on practical problems such as zip code recognition and speech synthesis a new wave of excitement spread across the artificial intelligence community. At that point, however, it was widely accepted that these models had no pretension of being biologically plausible

[8] For instance Laudan mentioned the early members of Royal Society of London who were convinced by mariners' tales of the existence of sea serpents and studied the properties of such serpents as an empirical problem [59].

[9] Note that what Bunge says about the relevance of practice in theory evaluation is here evoked even by Laudan but, in this case, from the side of science.

except of being interesting computational devices [77].

The evolution of neural networks clearly suggests that research hypotheses which fail to cope with more traditional questions of science (such as explanations or predictions) are not necessarily expected to miss other types of problems, such as those posed by technology. Moreover, for the reasons we previously seen on the overlapping activities of science and technology, they are neither expected to be exclusively engineering. Neural network, indeed, even after their exiting technological successes did not stop providing cues for other scientific investigations, for instance, on specific aspects of learning (see, e.g., works on regularizations networks or more recent studies on deep learning). But we dare say that this time the way in which neural networks approached science has been much more similar to the one presented by Bunge or Laudan and more in general by the disciplines which lies between science and technology. Furthermore, many other techniques may well exemplify the ideas we introduced before thanks to Laudan's and Bunge's contributions. In particular we think that pattern recognition research presents several elements of creativity, which might typically arise in various technological activities (e.g., in the form of operative theories), and, on the other hand, a high degree of pragmatic attitude towards more theoretical efforts (as the same way as that found in Laundan's problem solution).

In the field of pattern recognition there are several technical approaches that can be viewed as "theories of action" in the sense supposed by Bunge. Many techniques indeed arose from specific practical problems and created a real tradition over the years. A famous example is the case of kernel methods. They represent a family of techniques whose fundamental success relies on the so called "kernel trick". Basically kernels can be considered as a non-linear generalization of standard dot products. They allow one to handle input data, $(x_1, y_1), ..., (x_n, y_n) \in X \times Y$, as they were linearly-separable in the feature space $\mathcal{H}$ without computing directly the mapping $\Phi : X \to \mathcal{H}$ ("kernel trick"). By this technique one is allowed to solve a practical problem and in particular to deal with "virtual" representations of non-linear structures. Hence the kernel trick and more in general kernel methods work as "operative theories" in view of the fact they arose specifically to address crucial aspects of pattern recognition research (e.g., representation and generalization). Another interesting example of techniques which were introduced to deal with specific context of action is given by search algorithms. In both their blind or informed version (heuristics) such procedures are formulated just as a sequence of actions bearing one to an initial state to a goal state. We find that the searching strategies which they were built on (in the case of blind search, e.g., see the breadth-first search and the depth-first search) represent indeed a creative attitude of practical rationality described by Bunge.

In a sense it might be easy to define the field of pattern recognition as a problem-solving activity. But what does this mean in the light of Laudan's argumentation? It could mean, for instance, that a theory in pattern recognition can be used independently from the fact that it is true or well-corroborated. This indeed occurs several times in pattern recognition. Consider, for instance, the theories for object recognition, such as the shape-based or the appearance-based models (see [1] for a interesting historical account). Basically their employment was determined by the successes achieved in particular aspects of the problem (for instance shape-based models, like Biederman's theory, are suitable in dealing with the viewpoint invariance) not by their internal consistency and neither by their ability to account extensively the issue of object recognition. Similarly the application of evolutionary game theory in clustering problems occurs basically regardless of the ability of the theory in explaining pattern recognition mechanisms. Yet, the same happens Markov Chain Monte Carlo methods whose application differs with the fields and accordingly with the criteria adopted to evaluate the adequacy of a solution. Therefore while in the field of cognitive science such methods are asked to solve the gap between rational models of cognition and psychological processes [98], in the fields of pattern recognition and machine learning they are asked to cope with high-dimensional models and massive datasets[9].

1.7 Summary

In conclusion, with the contributions of philosophy of technology and philosophy of science, we argued that we should rethink the classical dichotomy between science and technology in order to understand why the field of pattern recognition could be properly considered as both science and engineering. Specifically we suggested that, on the one hand, engineering activities are often continuous to the work of ordinary science and on the other hand, many scientific enterprises can be characterized as an engineering effort. We described these interactions by means of some analytical tools developed by Bunge (in the philosophy of technology) and by Laudan (in the philosophy of science). We argued that Bunge's operative theories can contribute to shed light on the creative strength of practical rationality within the field pattern recognition as well as within other engineering disciplines. While Laudan's problem solving approach can in turn provides interesting cues to look at pattern recognition in scientific terms thereby exploring the logics and the pragmatics of problem solution.

2

Pattern recognition as a computational problem

In this chapter we would like to consider in which way problems of pattern recognition are formulated, pointing out some basic characterizations. Note that the work of pattern recognition has been often associated to the problem of induction which is a fundamental notion even in the field of philosophy. Therefore we will also attempt a brief philosophical account of such notion so as to highlight some interesting distinctions between a modern and an ancient conception.

In order to better understand what pattern recognition deals with we will devote the last part of our chapter to investigate the connections with some related areas. Specifically we will tray to outline how pattern recognition is related to artificial intelligence (is pattern recognition a sub-field of artificial intelligence) and to machine learning (do pattern recognition and machine learning coincide?).

2.1 Pattern recognition problems

Generally when we are dealing with a pattern recognition problem we would like to discover regularities among data. basically we start from a number of observations that can be represented into a feature vector space so that each object can be characterized in terms of a finite number of properties. Note that feature-vector representations are extremely attractive because geometric spaces provide powerful analytical as well as computational tools. However, besides the typical feature-based representation, we may describe objects even by structural and similarity information (e.g., by graphs or trees). [1]

Note that for most practical applications the selection of appropriate features is a crucial task as it reduces the variability within classes. To this end a common pattern recognition procedure requires a step for the extraction of “strategic” features (see the feature extraction step in figure 2.1). Once the features have been selected we need to find a “good” classifier, that is a function which collect similar objects and separate different ones. This task may also be called generalization and consist

[1] For a discussion on similarity-based representation see [79]

Figure 2.1: Typical steps in a pattern recognition problem.

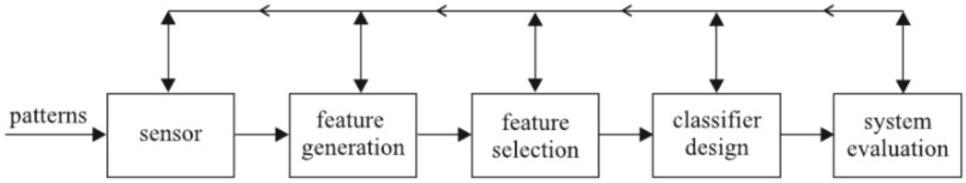

mainly in a learning activity because a pattern recognition system should learn somehow how to relate objects. Note that this task can be fulfilled in several ways. For instance, we could generalize the information provided by a training phase. In this case we talk about "supervised learning" because the system is presented by some exemplars, that is by a set of objects provided with their corresponding labels. After presenting training exemplars it follows a test phase, in which the system has to label the unseen objects. By contrast when there is no training step the system is asked to discover the "structure" underlying data without any supervision. For this reason we should talk about unsupervised learning or clustering. Then, between the supervised and the unsupervised condition there is an intermediate situation, often called "semi-supervised learning" in which classification occurs with both a small set of labelled data and a large number of unlabelled data. Another type of scenario is that described by reinforcement learning. In this case the system is supposed to discover optimal output interacting with the environment by a mechanism of trial and error. Finally a pattern recognition problem has to be evaluated. Validation may vary based on the learning scenario (for instance, where it is supervised or unsupervised). But usually the sense of validation is to estimate the error of the obtained classification and might be done by counting the misclassified test items (i.e., objects labelled during the test phase) [2].

2.2 Pattern recognition as induction

It is widely acknowledged that pattern recognition and machine learning provide computational tools to perform induction inferences [40, 122, 117]. Therefore, finding a pattern is often a synonym of induction. Harman and Kulkarni, indeed, stated that "an inductive method is a principle for finding a pattern in the data that can be used to classify new cases or to estimate the value of a function for new arguments. So the problem of finding a good inductive method is sometimes called a pattern recognition problem" [40, p. 22]. Yet, other researchers expressly advocated that machine learning is concerned with algorithmic induction [75]. Such correlations

[2] Note that the evaluation of a clustering problem could be much more critical since "there are a huge number of possibilities regarding what will be done with it" [118, p. 2]

were used even to show that machine learning and the philosophy of science are involved in similar tasks as both fields are concerned with the tasks of selecting and testing hypotheses [122].

In philosophy induction can be described as a process of generalization going from experience to the formulation of laws, from what is already known to what is yet unknown. Commonly is defined in opposition to the deductive reasoning. The latter is notoriously formulated by syllogism like "all people are mortal. I am a person. So, I am mortal" and it is associated to a sort of "perfect reliability" [40] which induction hardly provides. Indeed induction derives hypotheses from a finite number of observations so that its predictions are preferably expressed in probabilistic terms. This led many philosophers to address the problem of induction (Hume, for instance, was one of the most famous) posing the question whether one is rationally justify in drawing conclusions about future things or events. But even if it raised many skeptical reasonings the philosophical evolution of induction was inextricably tied to the development of one of the most fundamental form of knowledge, i.e. (modern) science.

2.2.1 An ancient and a modern sense of induction

In the modern age induction soon became the name to indicate the scientific method. Rather than being a simple matter of enumeration, induction outlined the methodical work of science which collects information and repeats experiments so as to infer causal relations and thereby formulate laws of nature. Accordingly, within modern philosophy the notion of induction were essentially associated to the problem of how justify scientific statement. Basically modern philosophy regimented the classical enumerative induction according to the criteria of the new science (consider, for instance, the influence of Bacon's method).

But modern induction is rather different from other early accounts and in particular from the Aristotelian conception. In Aristotle's writings we find at least two ways of thinking of induction. The former is essentially the enumerative version that was in time criticized by modern philosophy, while the latter coincides with the ability of grasping "universal forms" from individual experiences. Therefore the second Aristotelian meaning is strictly related to the overall process regarding the knowledge of universal concepts. In the Posterior Analytics Aristotle introduces induction as a form of abstraction which allows humans to pass from sensory information to the knowledge of things. This passage, hence, would consists in "transforming" sensory data into cognitive representations[3] which in time became the building-blocks of human cognition. In the end we may compare the process by which humans achieved knowledge " to the way in which order is restored in a battle after a rout.

[3]Here the notion of representation is rather different from that one could find in modern philosophy. According to Aristotle a cognitive representation is a sort of immaterial "presence" of the object within the soul and should not be confused with the (modern)representationalism criticized later by Wittgenstein.

First one man stops in his flight, then another, then one more, until there is a nucleus for real work. Similarly, the flow of fugitive impressions stops at one point; a similar impression comes along, is arrested by the first, and reinforces it; thus, after a time, there is formed a single experience. This supplies the starting point for the conscious process by which a system of conceptions is formed" [121, p. 155].

2.3 Pattern recognition and artificial intelligence

Pattern recognition is often considered a sub-area of artificial intelligence and, more expressly, as a field dealing with a number of typical problems and techniques sprung from a process of overspecialization. On the other hand, the idea that "pattern recognition encompasses a wide range of information processing problems of great practical significance, from speech recognition and the classification of handwritten characters, to fault detection in machinery and medical diagnosis" [13, p. 1] is rather popular in the literature. As a consequence, its development and its current role could be viewed on a par with many other sub-areas which finally have made progress on specific topics (machine learning, computer vision etc.) or methodological approaches (neural networks, support vector machines, etc.). But, there is also a different perspective which holds exactly the other way around and, hence, insists on the primacy of pattern recognition over the other areas of artificial intelligence. In this respect, pattern recognition is at the very core of cognition and, accordingly, of each attempt to build intelligent machines. Interestingly, more than 40 years ago, Hubert Dreyfus had already noticed that "the resolution of the difficulties which have arrested development in game playing, language translation, and problem solving presupposes success in the field of pattern recognition." [24, p. 97] So, how should we figure the relationship between pattern recognition and artificial intelligence? how do these fields really stand to each other?

From an historical perspective, it is undeniable that pattern recognition and artificial intelligence are intertwined and at the early stages of their development the two really overlap. As Duin pointed out [26], before artificial intelligence were officially established at the famous workshop in Dartmouth, in 1956, there was no formal distinction between pattern recognition and the other research projects dealing with the mechanization of the mind. Indeed, at the same Western Joint Computer Conference (WJCC55) held in 1955, we can find a session on machine learning which included some papers explicitly devoted to the problem of pattern recognition. However, it is interesting to note that all presented works were ultimately perceived as mutually related and, in particular, as reflecting distinct levels of increasing (biological) complexity [4].

[4]Indeed, Willis Whare introduced the papers of this session pointing out that "these levels of complexity may be likened respectively to the initial organizational efforts of neural nets in the earliest days of evolution (Farley-Clark), the learning of pattern recognition by a maturing individual (Selfridge-Dineen), and the action of an entire man as he copes with the complicated

Along the way the computational approach to pattern recognition became more and more defined and application-oriented (see, e.g., the abundance, at that time, of models for recognizing symbols in written texts). This process of consolidation involved researchers from different domains including, among others, psychologists, engineers and computer scientists as well as physics and mathematicians. During 70's the separation between pattern recognition and artificial intelligence went evident and culminated with the institution of separate journals and conferences. In 1968 it was founded the first journal on pattern recognition and two years later one which was completely dedicated to AI. Most importantly, as Duin put it "after the first International Joint Conference on PR, Washington DC, 1973, it was decided that in the even years there would be an international conference on PR to avoid a collision with the international conferences on AI, which were organized in the odd years from 1969." [26] Therefore, in the light of its historical evolution, pattern recognition may be reasonably considered a sub-field of artificial intelligence as a by-product of a process of fragmentation. From this standpoint the establishment of particular resources for the growth of research, along with the connected social effects (e.g., "narrow" membership, stereotyped research, etc.), has for sure contributed to the idea that pattern recognition is in a way "subordinated" to artificial intelligence, even though due mostly to incidental reasons. Nevertheless, this did not prevent the two communities to support joint activities (see, e.g., journals covering both domains like International Journal of Pattern Recognition and Artificial Intelligence or IEEE Transactions on Pattern Analysis and Machine Intelligence).

Admittedly, according to Duin the separation between pattern recognition and artificial intelligence was a process which in fact expressed two distinct problem-solving approaches, equally important (i.e., without subordination) but going in two opposite directions. Yet, Duin holds that these two areas diverge in a very scientific way and outlines such a distinction thanks to the popular comparison between two of the most influential philosophers of the Western culture, Plato and his pupil Aristotle. Hence, artificial intelligence, following a "top-down" approach, typically found in Plato's doctrines, would move from some fixed, given concepts to the objects of real world [5], whereas pattern recognition would proceed in the opposite direction, generalizing from observations which therefore come first, as it is supposed to happened in a genuine "bottom-up" scheme.

However, if we want to characterize the fields of pattern recognition and artificial intelligence from a scientific standpoint, that is looking at the sort of phenomena these two areas get in touch, we will suggest further considerations. More than a different working style, what relates pattern recognition to artificial intelligence is a sort of profound dependency precisely as the one connecting categorization to all

problems of his life (Newell)." [119, p. 85]

[5] Duin calls this process also "adaptation" [31, 30]. Note that he associates this approach also to the troubled story of neural networks and the related incentive towards alternative researches based on logic and reasoning.

intellectual activities [6]. Pattern recognition, as we have seen before, has been indeed understood as a way to point out categorization or, more basically, to "see the one in many" [120]. More recently, it has been explicitly argue that in the analysis of cognition everything, in some way or another, can be reduced to pattern recognition [61]. And even in the philosophy of science it has been proposed that "we use our somewhat innate capabilities of pattern-recognition to find our way through the maze of facts and hypotheses" [109, p. 31]. Therefore, as a cross-cutting notion, pattern recognition looks more like a starting point or a principle of discovery than a side effect of the development of artificial intelligence.

2.4 Pattern recognition and machine learning

More than ten years ago a paper by Duin, Roli and de Ridder noticed that "although many pattern recognition scientists seem pretty sure about the cultural identity of their research field, [...] the recent developments of closely related disciplines (e.g. machine learning and neural networks) and the increasing number of research issues that pattern recognition shares with such disciplines"[29, p. 1] makes difficult to answer questions about the cultural identity of pattern recognition[7]. The complexity of this scenario then increases when we consider that pattern recognition and machine learning are in fact considered as "two facets of the same field" [12, p. 1]. From a practical point of view, this intuition often results in a general attitude towards exchangeability, that is in presenting several pattern recognition works on the "machine learning" platform and vice versa [27]. But if we look at some traditional features which generally characterizes the establishment of a research community (conferences, journals, etc.) we will conclude that pattern recognition and machine learning are two distinct areas since both have their own journals and conferences. Moreover we would genuinely expect some differences from two areas which were designated under two distinct terms. One the one hand such differences may be explained just in terms of historical or sociological factors, maybe associated to the opposition between influential traditions (e.g., the symbolic and the connectionist approaches). On the other hand there might be the possibility to identify more profound difference. This is, for instance, the idea suggested by Duin in an informal discussion where he said that, in fact, between machine learning and pattern recognition "there in an essential difference in focus" [27]. According to Duin pattern recognition and machine learning would diverge because the former is about the recognition of patterns itself, in general, without restrictions, whereas the latter is

[6]This is not only a basic tenet of modern cognitive sciences [17] but also a cornerstone of any philosophical theory of knowledge since antiquity. Indeed, how should we understand the debate which traversed the whole history of philosophy on the constituents of thought but for a testimony of the crucial role of categorization?

[7]Actually that paper was devoted to statistical pattern recognition but we think that their motivating questions could be suitable even for pattern recognition in general

primarily concerned with the study of learning process. We do not know how much this theoretical distinction affect the concrete practice within the two communities but we think that the idea of distinguishing the recognition from the learning has some reasons.

The fact that pattern recognition is strictly correlated to the process of learning was known since the early days of artificial intelligence. In 1955 Selfridge, for instance, at the western joint computer conference claimed that pattern recognition "leads naturally to studying other processes, such as learning" [103, p. 91]. Its definition, however, seems to not really require a reference to learning and in fact even the development of pattern recognition techniques may occur without learning (nearest-neighbour methods come immediately to mind). In this respect pattern recognition may be referred to representation and, from a philosophical standpoint to one of the Aristotelian ways to conceive induction. As we suggested above Aristotle meant induction even as a form of immediate abstract "representation" of what one can perceive by senses. Interestingly, in a purely Aristotelian formulation we would not say that pattern recognition is a real process as it seems conceived outside the logics of time and it looks like a type"intuition". Therefore pattern recognition could be more properly associated to the Aristotelian idea of induction as "abstraction". While machine learning could be more directly connected to the modern account of inductive inference. Interestingly even a more recent study on the relationship between the problem of induction and pattern recognition acknowledged that "rules of classification must be carefully distinguished from inductive methods for finding such rules". And while rules of classification are more concerned with categorization or function estimation the induction method is about the use and the choice of such rules. therefore, in very informal terms we could say that when the mind "meets" the world many ingredients for categorizations are present but one needs experience and training to understand how they can be "used". With respect to the characterization of the field of pattern recognition this differentiation may in turn may lead to a "list" of primary issues, such as pre-processing procedures (how to represent objects, how to choose proper features etc.) [27].

2.5 Summary

In this chapter we sketched out some classical components of a pattern recognition problem. Then we presented pattern recognition as an inductive problem. We observed that, besides the modern formulation (which is probably the most popular) there is also another important sense of induction which arises from Aristotle's philosophy. The two meanings in the end could be useful to distinguish between learning (modern induction) and recognition (Aristotelian induction as abstraction). These elements then could in turn be used some special focuses within machine learning and pattern recognition, two areas which are often seen in strict correlation (and sometimes really as the same thing). To better characterize the role of pattern recognition we also considered some connections to the broad area of artificial intelligence. In this respect we may conclude that the two positions sketched above - i.e. the one asserting the primacy of artificial intelligence and the other maintaining the priority of pattern recognition - are in fact partly true. For we have good reasons to believe that historically the field of pattern recognition has come out of artificial intelligence as a particular scientific community, with its own journal and conferences. But, on the other and, we are also provided with several convincing arguments which make it clear that pattern recognition is the very building block of intelligence and, accordingly, of its mechanical version.

3

Pattern recognition as categorization

Categorization is another way to paraphrase the subject matter of pattern recognition. When we speak about categorization we usually suppose a situation which bears a strong resemblance to a classical pattern recognition task. Indeed when we are involved in a categorization task basically we are asked to place particular objects into classes (that may be called "concepts" or "category"). But this problem is a central issue even in the field of cognitive psychology, as well as in the area of pattern recognition.

Cognitive psychology is concerned with the study of several mental processes such as perception, memory, attention, etc. Its approach is based on the assumption that the brain can be in principle described as an information-processing system. Therefore analogies to computer science are very common as it is believed that, like a computer, the brain takes information from the world, it converts the received data into some representation and then it produces a (sensible) result. This analogy is particularly significant in the study of categorization. In this respect, cognitive psychology has developed several computational models that could be easily associated to pattern recognition techniques. Indeed, many such models result in algorithms which produce quantitative predictions on specific categorization tasks. But the similarities between pattern recognition and cognitive psychology also have deep historical roots since many researchers have contributed significantly in both camps (Marr or Biedermann are two well-known examples). Nowadays such connections are stressed by interdisciplinary projects which stem from the practice of cognitive sciences and aim at a closer link to machine learning.

But behind the apparent formal analogies, pattern recognition and cognitive psychology hide two different perspectives on categorization. Looking at the research activity of cognitive psychologists, for instance, we may have the (strange) impression, on the one hand, to deal with abstract data sets that have few relations to the every-day classification tasks (e.g., sorting images, ordering objects in the desk, etc.) and, on the other hand, when we look at the field of pattern recognition, by contrast, we may come to the conclusion that we handle useful tools which are far away from the study of cognitive phenomena. Therefore, we might realize that if we face the issue of categorization from the standpoint of one of these two particular fields we will not be guaranteed to account the same phenomena. Rather, working separately in one of these two areas we may end up with an original conception of what categorization is and how it should be studied. Therefore, it makes sense to pose some questions like: do cognitive psychology and pattern recognition really address the same problem? what common attitudes do they share? and, on the other hand, what peculiarities do they exhibit? To what extent are their methods effective in performing categorization?

Recalling Laudan's terminology we would say that we are interested in how pattern recognition and cognitive psychology shape their problem-solving activity. Specifically we would like to address a particular instance of categorization activity, that is unsupervised classification. The latter may fall under different labels (e.g., unsupervised learning, unconstrained classification, clustering, spontaneous categorization) and in general corresponds to the task of grouping a number of objects without receiving any previous training or guidelines. Humans are very familiar with such tasks (a lot of every-day experiences come to mind, e.g., naming things, arranging clothes in the closet, etc.) whose intrinsic spontaneity in fact constitute the main challenge of computational approaches.

In our discussion as well as proposing some general considerations about how cognitive psychology and pattern recognition conduct their modelling activity we will develop a concrete comparison between some models and methods coming from both fields. Note that the selection of such models has not been motivated by particular experimental goals but for the interest in making our discussion more effective. We do not have any pretension of exhaustiveness and we hope to offer just a contribution to the interdisciplinary dialogue on these neighbouring areas.

3.1 Formal approaches to categorization

Categorization is probably one of the most intriguing problems in the history of philosophy. It has been studied from several points of view and in different philosophical traditions (e.g., continental, analytical, pragmatics, etc.). Nowadays, categorization may be approached in a different way, by means of formal models and techniques allowing quantitative predictions. This approach is greatly present in both cognitive psychology and pattern recognition and guides the very development of computational accounts of categorization.

In order to point out how cognitive psychology and pattern recognition conceive spontaneous categorization we provide a brief outline of their common methodological inclination. At the moment we prefer to leave aside the question whether pattern recognition and cognitive psychology produce models like those of traditional science because such a problem would deserve a discussion apart [1]. But in reference to our discussion we could take for granted that a "formal model is one that unambiguously specifies transformations from one or more independent variables to one or

[1]For a general overview on the role of models in science see [36].

more dependent variables. In the case of formal models of categorization, one independent variable is category structure, and one dependent variable is categorization accuracy" [123, p. 102].

From a general point of view models in science can be used for two main purposes. They can be built to represent a selected part of the world [2] (what is usually called "target system") or to interpret the laws and the axioms of a theory (e.g. Euclidean geometry). In a way the problem-solving activity of pattern recognition and cognitive psychology could be understood as a modelling activity. Indeed, ultimately, both aim at formalizing a type of inference which takes some data as input and returns a grouping hypothesis as output. With respect to the great variety of models (analogical models, computational models, explanatory models, heuristic models, etc.), what is common in pattern recognition and cognitive psychology is the tendency to express that inference in a formal way, that is without ambiguities or vague concepts, and with the support of computers. To this end, formal models of categorization rely on a vast collection of mathematical tools with the advantage of getting quantitative results which can be easily compared. In addition, they can count on the progress of computer science and the great expansion of the algorithmic language. In particular these elements make models suitable for computer simulations, which seems almost to play the role of definition in the traditional philosophy. That is, instead of explaining categorization by addressing the typical philosophical question about the essence (what categorization is?), cognitive psychology and pattern recognition moved their efforts towards an operational approach, looking at the process and the aspects of measurement (how does categorization work?).

In order to understand this modelling process we would like to stress two fundamental components: the elaboration of assumptions or hypothesis about the mechanism of categorization and the formulation of validation methods for the assessment of model predictions. We think that these two aspects can provide a common ground for our comparison, besides the particularities of each field (e.g. terminology, technical details, etc...).

- Assumptions. One of the possible ways to approach a model of categorization is to analyse the set of implicit or explicit ideas, which influences the formalization of the proposed method. Sometimes this phase doesn't constitute a distinct or an expressed passage, however, even if implicitly, assumptions may constrain the technical choices in a very deep way.

- Validation. This component provides the principles that are used in evaluating algorithm performance and comparing results of different models. As well as being a practical tool, methods of validation offer an indirect measure of coherence. Indeed, by looking at the validation criteria one indirectly express an ideal picture of what categorization should look like.

[2]Note that "depending on the nature of the target, such models are either models of phenomena or models of data" [36]. For a discussion on models of data see: [111]

In the next two sections we will briefly introduce the approaches of cognitive psychology and machine learning keeping in mind these two aspects (assumption - validation) so as to draw some general considerations.

3.2 Spontaneous categorization in cognitive psychology

In cognitive psychology a simple way to understand the process of spontaneous categorization is to evoke a typical experimental scenario: "participants are asked to divide some stimuli into categories that are intuitive, without any corrective feedback" [87, p. 1]. This is the behaviour that cognitive psychology studies and tends to formalize in models of unconstrained categorization. Note that the objective in this context does not consist in solving a specific classification task but in examining and representing a human phenomenon. Consequently, cognitive psychology is more inclined to express accurate assumptions about categorization, relating models to some well-defined theory (such as the exemplar or the prototype theory, or the Bayesian framework, etc.) or claiming which principles inspire the technical choices. For example, there are some models supporting the hypothesis that categories spring from similarities among objects and then develop methods exploiting such information. At the same time, others suggest that a "category is coherent if it fits well with our overall understanding of the world" arguing "that explanations based on similarity are inadequate" [87, p. 1]. An alternative view states that the process of categorization serves certain functions of the organism and puts forward models based on category utility. An example of this approach is given by methods addressing questions such as: how useful is the category in predicting the features of its members? or which category is more appropriate to communicate a collection of properties? [20].

From a philosophical point of view the instantiation of these theoretical coordinates provides explanations and arguments on the nature of relations joining objects together, or, in other words, the reasons underlying grouping mechanism. Some people, for instance, argue that the formation of general entities (group or categories) is a matter of general knowledge since that each concept is inseparable from the overall knowledge of the world and, in turn, without this knowledge it seems difficult to appreciate the significance of a concept [70]. Others support the idea that objects are grouped because of their mutual similarities (see, for instance, the models based on multidimensional scaling techniques). While all these considerations put forward mainly qualitative information and support the general setting of the models, mechanisms of validation provide quantitative tools and represent a complementary way to speak about the structure of a category (e.g., the degree of coherence or the requirements for a good category).

Usually, the evaluation task occurs with relatively small artificial data sets (points or abstract figures) both with humans and algorithms. But such data sets are often the result of a precise experiment design expressing a number of properties which may be relevant for the model testing. The main purpose of evaluation is to assess the ability of models to "minimize the quantitative difference between their outputs and some empirical observations" [123, p. 110]. Such observations commonly are collected while a sample of humans performs a (categorization) task during appropriate experiment sessions. A step further in this evaluating process could be the comparison with different model predictions. Indeed there might be models which provide a better account of empirical phenomena and which should be preferred to other competitors. Note that there are several options to evaluate the performance of models. Many approaches are based on the sum of squared errors (SSE), that is the distance between the output of models and the human predictions. However, such a measure may not be always adequate. Indeed Wills and Pothos pointed out that "two models can have the same SSE in cases where most theorists would agree one is superior, and better SSE can sometimes indicate a less adequate model" [123, p. 111]. An alternative view suggests that model evaluation should include more considerations for the qualitative properties of model behaviour. According to this view, the evaluation should be arranged against a criterion of ordinal adequacy so as to assess wether models capture the ordinal properties of a data set. For example, a criterion of ordinal adequacy has been included in some recent studies on category intuitiveness [87, 88]. Such studies try to assess how good are models in predicting intuitive category structures. Interestingly, the proposal leads models to the very heart of unsupervised classification providing a measure for intuitiveness, or, in other words, for spontaneity. Note that the problem faced by this investigation is related to many ordinary experiences: "why do we consider a category like 'chair' as intuitive (coherent), a category like 'games' as less intuitive (in the sense that people disagree more about the membership of this category), and a category composed of 'babies, the moon, and rulers' completely nonsensical?" [88, p. 84]. To better understand how intuitiveness is conceived consider, for instance, the figure 3.1. Assuming that such images represent a psychological space where each point is taken to correspond to a physical stimulus, we may find out that classification in top graph is perceived as more intuitive than classification in the bottom graph, maybe, because of the difference between the ratio within-between similarity. In the studies, then, intuitiveness has been measured by the frequency of the preferred classification (i.e., that more produced by humans) for each data set, that is measuring how much participant agree in producing classifications (thus a strong agreement would indicate a high level of intuitiveness while a large disagreement would be a signal of low intuitiveness). Interestingly, as well as proposing a dependent variable for unsupervised classification (i.e., category intuitiveness), this study evaluates models by an ordinal criterion. Indeed, it assesses models not really to measure the error with respect to each data set but to compare the order of (preferred) classifications produced by models and humans. Note that in this way it does not care to find the right classification for a single stimulus set, but rather to capture the degree of

Figure 3.1: Items in psychological space. The top graph represents an intuitive category structure whereas the bottom graph a corresponding less intuitive one. Reproduced from [87] with permission.

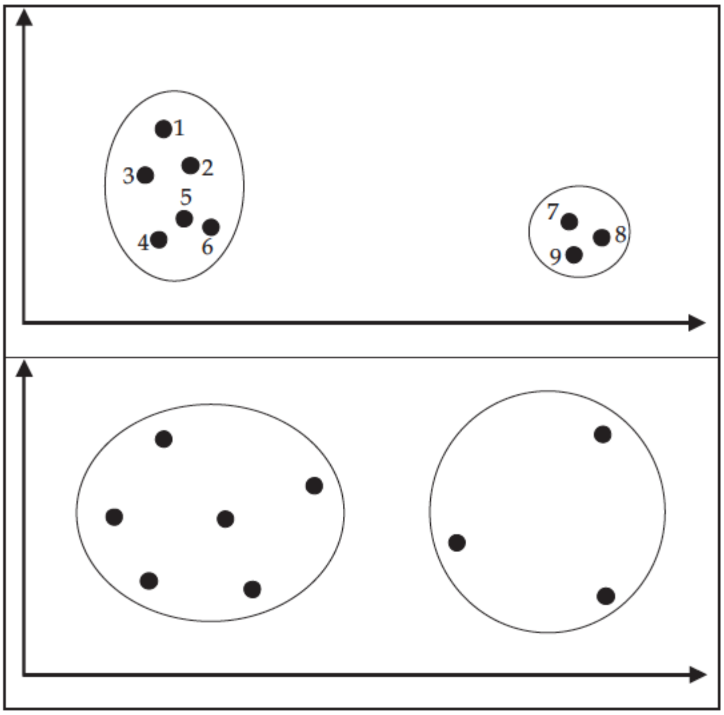

intuitiveness of different category structures. We will return on this basic test later on, since such studies deal with a crucial notion for unsupervised categorization.

3.3 Spontaneous categorization in pattern recognition

Spontaneous categorization somehow corresponds to what in pattern recognition and machine learning is called "clustering" [117]. However this problem does not belong exclusively to the domain of pattern recognition since there are many other research areas which are interested in discovering regularities among data (e.g., data mining, statistics, etc.). Hence, when we speak about clustering we are in fact considering a cross section in computer science. However, the problem of clustering is historically associated with the dawn of artificial intelligence, when computer scientists started studying and designing intelligent agents. The key elements of machine learning and pattern recognition (representation, generalization, evaluation) emerges as a special instantiation of human skills and at the same time as an alternative way of exploration [3]. Thus, at the very beginning of clustering there was the attempt to make machine intelligent and, possibly, to bring computers near to human qualities. But, over the years, these origins have been disregarded in favour of a greater emphasis on application-oriented approaches. As a result, in computer science spontaneous categorization has lost the "genuine" reference to the agent of categorization preferring a more abstract characterization. Indeed it is usually supposed that: "the field of machine learning does not study the process of learning in living organisms, but instead studies the process of learning in the abstract" [117, p. 652]. In clustering research the problem is to discover some "structures" on the underlying space of instances and, in so doing, to provide effective solutions to many ordinary situations. For instance "an online retailer might want to cluster his customers based on shopping profiles. He collects all kinds of potentially meaningful information about his customers [...] and then wants to discover groups of customers with similar behaviour. [...] It is not specified beforehand which customer should belong to which group it is the task of the clustering algorithm to work that out" [117, p. 652].

Focusing primarily on problem solution, the field of pattern recognition has developed a different kind of assumptions compared to cognitive psychology. Generally, the majority of clustering algorithms are constrained by computational principles (complexity, efficiency, robustness, reliability, applicability, etc.) or by general properties that a solution ought to satisfy (see, e.g., the axioms proposed by Kleinberg: consistency, richness, scale-invariance [51]). In the last few decades the computa-

[3] Usually introductions to classification and clustering stress this origin: "Classification is both an ancient discipline (Aristotles classification of animals, plants, and other objects is still largely valid), and a modern one. Classification is a way of thinking about things, rather than a study of things in themselves, and so it draws its theory and applications from the complete range of human thought and experience" [41, p. 1].

tional and the engineering concerns prevailed against other factors (e.g., psychological o philosophical). However the development of clustering techniques implies at least an intuitive idea about what a categorization mechanism ought to be. It may remain undisclosed or just outlined but such an idea puts forward an hypothesis on how categorization works and after all influences the entire design process. A widespread supposition (that we can call "feature-based approach"), for instance, is the idea that categories are defined on the basis of a collection of properties or features. From this point of view, instances are represented by a collection of features and grouped in accordance with the properties required by categories. In terms of coherence, the notion of a category coincides with a set of attributes and induces fixed boundaries in the same way as the set-theoretical model introduced by Lakoff [4]. Similar considerations emerge in the analysis of validation methods which reveal different points of view.

Validation in clustering occurs in several forms. A recent review of common methods has pointed out some popular practices different types of data-sets (artificial, classification benchmark or real world). But in all cases, there exists a source (e.g.: expert people or particular assumptions on the data generating process) providing a sort of ground truth against which the algorithm performance can be compared. Another possibility is given by the so called "internal clustering quality scores" approach. In such scenario, one can use many different measures (e.g.: sum of square distances to cluster centres, ratio between-within similarities, likelihood scores, etc.) to express the goodness of clustering and consequently to determine the internal cluster coherence[5]. However, there is no agreement on the reliability of these evaluation procedures and different judgments are possible. Some people advocates for a "universal" way to compare algorithms legitimizing the idea of a single right solution for clustering, independently from contextual information. Others take this further and maintain that the right answer is determinable by the data (alone, without reference to the intended use): "the data should vote for their preferred model type and model complexity" [15, p. 1]. An alternative view shifts the point of interest from data to the specific purposes that clustering is asked to serve. According to this perspective: "clustering should not be treated as an application-independent mathematical problem, but should always be studied in the context of its end-use" [118, p. 1].

[4]The set-theoretical model lies at the heart of the classical theory of categories according to which "given any property (or collection of properties), there exists a category in the world consisting of the entities that have that property" [56, p. 159]

[5]It is worth mentioning that in computer vision there is a benchmark (the Berkeley segmentation benchmark) which actually provides an empirical basis for research.

3.4 Points of interaction

In the light of these few observations now we can derive some points of interactions. Indeed we may say that in the end cognitive psychology and pattern recognition present a specific view with respect the objects, the process and the theory of categorization.

- *Object of categorization.* This point mainly regards the type of data sets used. How do cognitive psychology and pattern recognition look at the object of categorization? From the brief account provided above we can observe two distinct attitudes. Most of the time in cognitive psychology we cope with artificial data sets whose final usage is restricted to the psychological inquiry. As the example in figure 3.1, they might be very abstract, simply representing some stimuli in a psychological space, and with a small number of items (e.g.: 20-30 objects). On the contrary, pattern recognition, as well as managing artificial data sets, handles "real-world" objects (e.g.: images, texts, biological data, etc.), whose size normally exceeds hundreds of items.

- *Process of categorization.* In both disciplines there is the idea of modelling a process. But what kind of process is the field of pattern recognition looking at? In the majority of cases categorization tasks are primarily a matter of engineering, concrete problems which need effective solutions. This emerges clearly in the evaluation practice where algorithm output is compared against a ground truth which makes sense with regard to statistical information (i.e. the process generating data) or the knowledge of experts but which is almost indifferent with regard to human thinking. By contrast, cognitive psychology is interested in categorization as a human activity and builds its models on the correlation between human response (the empirical observation) and the output of algorithms. These distinct characters arise in the elaboration of preliminary hypotheses as well. In fact we observed that models of cognitive psychology are frequently inspired by theories of categorization whereas pattern recognition algorithms express their hypothesis in an implicit way. Nevertheless, it is worth noticing two facts. On the one hand, the lack of a specific interest in human categorization does not prevent pattern recognition algorithms from advancing useful applications in many real world categorization tasks. On the other hand, because of their explanatory function, psychological models of categorization are not necessarily asked to be computationally efficient or robust (even if there are some exceptions such as the models based on optimal inference, see [98]) but, rather, to be as close as possible to human preferences.

- *Theory of categorization.* What kind of modelling activity can we draw from cognitive psychology and pattern recognition? Which relation between the practical level and the theoretical one? The main work in the community of pattern recognition is not devoted to the development of a theory of categorization, at least in the sense implied by ordinary science [112]. The theories developed by pattern recognition are in principle technological [16]. They work indeed as ordinary science (e.g., they are built on idealization of the problem, they employs a great deal of mathematical tools, they undergo several tests) but are constrained by applicative issues. So, we could also say that they are theory for "applications". For what concerns cognitive psychology we observe that the theoretical work is strictly joined to the amount of intuitions about categorization. Generally models of categorization do not live in solitude but mostly in families of related models. Actually this may hold even for pattern recognition but in this case the networks of methods (e.g., support vector machines) are based on common technical approaches whereas models of cognitive psychology are usually related by common explanatory hypotheses. Such models are built to express in a formal language what has been articulated in an an informal way (for instance, by argumentation and observations). Then, experimental results (derived from comparing model predictions against empirical observations) are used to get information about the explanatory hypothesis underlying models and their practical relevance goes beyond the logic of problem solution.

3.5 A practical comparison

In the next sections we will extend the comparison between cognitive psychology and pattern recognition to a testing phase. We did not undertake such a comparison as a real experimental activity, whose accomplishment has some traditional requirements [6], but essentially as a concrete opportunity to make practice with the methodological attitude of both fields. In fact, the main purpose in this phase is to observe how algorithms behave in a different context from their usual application without any pretension of completeness. Specifically we would like to know how models of cognitive psychology deal with data sets commonly used in the field of pattern recognition, and, on the other hand, how some pattern recognition methods cope with an experimental scenario developed for cognitive purposes.

From the field of pattern recognition we decided to select the techniques of k-means [46] and dominant sets [69]. While, from the field of cognitive psychology we opted for the generalized context model (GCM) [73] and the rational model [8]. A brief description of these models and techniques are presented below.

[6] To explore the relationship between experimental activity and computer science see, e.g., [6, 101]

3.5.1 The generalized context model

The generalized context model is based on the assumption that "people represent categories by storing individual exemplars (or examples) in memory, and classify objects based on their similarity to these sorted exemplars" [89, p. 18]. In the standard version of the model exemplars are represented as points in a multidimensional psychological space and the similarity between points is the decreasing function of their distance. Note that similarity is a notion highly context-dependent since it is modelled in terms of a set of selective-attention weights. This means that not all dimensions are always relevant for classification, i.e. the weight can stretch the psychological space along attended dimensions. Practically the model presents and initial training phase, in which the learner is presented with some exemplars of each category, and a test phase, in which both exemplars and new items might be presented. During classification the learner assigns the test items to one of the K_n available categories. Specifically new assignments are computed in this way:

$$P(\mathcal{F}|i) = \frac{b_{\mathcal{F}}[\sum_{j=1}^{n} V_{j\mathcal{F}} S_{ij}]^{\gamma}}{\sum_K b_K [\sum_{z=1}^{n} V_{zK} S_{iz}]^{\gamma}}$$

where $P(\mathcal{F}|i)$ is the probability that a new item i is assigned to the category F; $b_{\mathcal{F}}$ is the response-bias for category $\mathcal{F}$; $V_{j\mathcal{F}}$ denotes the memory strength of exemplar j with respect to the category $\mathcal{F}$; and S is the similarity between the exemplar j and the new item i. Finally, γ is a response-scaling parameter which affects the degree of determinism in label assignment. Hence, the probability of assigning an item to a particularly category is proportional to the overall similarity between that item and the category exemplars. The similarity between item i and exemplar j is given by

$$s_{ij} = e^{-cd_{ii}^{p}}$$

where c is a free sensitivity parameter which reflects the rate at which similarity declines with distance, while *determines* the shape of the function relating similarity to distance (e.g., exponential or gaussian). In the standard version of the model, the distance between i and j are given by the weighted Minkowski power model:

$$d_{ij} = \left[\sum_{m=1}^{M} w_m |x_{im} - x_{jm}|^r]^{1/r} \right]$$

where r determines the form of the distance metric (e.g., euclidean, city-block) and w is the attention-weight parameter referred to dimension m and with $0 \leq w_m \leq 1$ and $\sum w_m = 1$.

Note that in order to measure how good the model is in predicting category intuitiveness Pothos et al. provided an unsupervised version of the model (UGCM). In so doing they estimated the intuitiveness of a classification by considering how well each stimulus is predictable given the assignment of the other stimuli to their intended categories. "Suppose we are interested in evaluating a classification for a set of stimuli, {123}{456789} (the numbers are stimulus ids). We can consider each item in turn as a test item whose classification is to be predicted, and all the other items as training items whose classification is given" [88, p. 88]. The model parameters are adjusted until the predicted classification probabilities for individual test items are as close as possible to the classification of interest.

The unsupervised version of the model suffers from a combinatorial explosion when the number of items increases (indeed all assignments have to be explored until the good one is found). So in order to employ the model on larger data sets we adopted a semi-supervised version (i.e., we presented the learner with one randomly selected exemplar per group and the rest of unlabelled items). Our adaptation is based on the "translation" of some variables:

- The bias for category $b_{\mathcal{F}}$ = prior probability associated to category $\mathcal{F}$, $P(\mathcal{F})$
- The memory strength of exemplar j with respect to category $\mathcal{F}$, $V_{j\mathcal{F}}$ = the posterior probability $P(\mathcal{F}|j)$

Substituting these variables we obtain an iterative method which computes the posterior probability $P(\mathcal{F}|i)$ on the basis of the previous step. The initial probability values can be uniformly distributed or randomly selected. The only condition we impose is that for each item the sum of its probability distribution (over categories) is equal to 1.

3.5.2 The rational model

John Anderson's rational model is arguably one of the most influential model of unsupervised categorization. Instead of focusing on the psychological processes involved in category learning, the rational model aims to explain categorization as an optimal solution to the computational problem faced by the human agent. Specifically, the rational model is based on the idea that categorization reflects the goal of optimally predicting the unseen features of objects, that is, we wish to be able to predict $P_i(j|F_n)$ the probability that (as yet unseen) dimension i of the object possesses the value j, given the feature structure F_n observed so far. Categories are formed to assist this goal. Hence, objects are assigned to categories in such a way as to make the feature structures of those objects most probable. As a Bayesian model, the rational model assigns a new object to the most probable category k given the features observed, $P(k|F)$. This posterior probability is based on the prior probability of category k, and the likelihood term, $P(F|k)$, the conditional probability of object features F,given membership in category k. Thus the rational model, like Bayesian models generally, makes it necessary to choose a particular prior concerning category membership. Anderson specifies this prior in the following way:

$$P(k) = \frac{cn_k}{(1-c)+cn}$$

where n_k is the number of objects assigned to category k thus far, n is the total number of classified objects and c is the so-called coupling parameter. This parameter governs the probability that a new instance will receive an entirely new label. In other words, the coupling parameter determines how readily new categories will be formed: for high values of the coupling parameter, larger clusters are favoured by the prior, whereas for low values the model will favour greater numbers of smaller categories.For what concerns the conditional probability, $P(F|k)$, we need to compute the probability of displaying features F on several dimensions given that F belongs to k: $P(F|k) = \prod_i f_i(i|k)$ (assuming independence between values on different dimensions). In the case of continuous dimensions, which is the case we will consider in this paper, the function of conditional probability takes the form of the Student's t distribution, with a_i degrees of freedom.

The original version of Anderson's rational model uses an incremental algorithm to assign objects to categories. Objects arrive sequentially for classification, and at each point in time new objects are assigned to the most likely cluster; though new clusters can be added, there is no other opportunity for re-partitioning the item set. Hence the algorithm provides only an approximation to the optimal estimates, and, given different presentation orders, can yield different partitions. Following Sanborn et al. we will refer to this algorithm as the local MAP (as in maximum a posteriori) algorithm [98]. Approximation is necessitated by the combinatorial explosion associated with considering all possible clusters, and, it allowed Anderson to test the model on data sets of several hundred items that would otherwise have been beyond reach.

Recently, Sanborn et al. proposed two alternative methods of approximation (Gibbs sampling and particle filters) that better approximate the posterior distribution and on the tests conducted fit human data at least as well as Anderson original (1991) algorithm [98]. They thus argued that these algorithms have greater psychological plausibility and may, more generally, provide a useful basis for rational process models. In their model evaluation, Sanborn et al. focused on particle filters, so the research reported here was aimed at taking a closer look at Gibbs sampling.

Gibbs sampling is a Markov chain Monte Carlo method that is widely used in approximating the expected value of a (continuous) function. As a Monte Carlo method, it uses repeated sampling as a means of deriving an estimate. Given a distribution $p(z) = p(z_1, z_2, z_M)$ from which we want to sample, we approximate the function by picking up an ideally infinite number of samples:

$$E_{p(z)}\left[f(z)\right] = \lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} f(z^t)$$

Each sample z_i is drawn from the probability distribution conditioned on the values of other remaining variables: $p(z_i|z_{\setminus i})$. Then, the sampling cycle is defined as follows:

$$z_i^{(t+1)} \sim p(Z_i | z_1^{(t+1)}, ..., z_{i-1}^{(t+1)}, z_{i+1}^{(t)} ..., z_k^{(t)})$$

With respect to the problem of categorization, Gibbs sampling is supposed to converge to the probability distribution over categories given a particular stimulus set. In this setting the state space of the Markov chain corresponds to the set of partitions and "transitions between states are produced by sampling the cluster assignment of each stimulus from its conditional distribution, given the current assignments of all other stimuli" [98, p.1152]. Note that even if Gibbs sampling converges to the desired distribution, it is standard to discard the early iterations, known as the burn-in samples, because they are not guaranteed to come from that distribution as they are biased toward the initial values. Likewise, the nature of thealgorithm means that there will be a strong dependency between adjacent samples. Thus in order to obtain (reasonably) independent samples the usual solution is to keep every n th sample, rejecting the rest.

3.5.3 K-means

K-means is one of the most popular techniques in unsupervised learning problems. In literature it is recognized as a partitional algorithm, that is as a procedure determining a partition of the items into k groups without providing a nested sequence of clusters. In the k-means algorithm the number of clusters is fixed a priori and all items are assigned to the cluster with the nearest mean. Given a set of observations $(x_1, x_2, ..., x_n)$, where each observation corresponds to a d-dimensional vector, the algorithm divides observations into k partitions $(S_1, S_2, ..., S_k)$ minimizing the sum of the within cluster distance:

$$\sum_{i=1}^{k} \sum_{x_j \in S_i} \|x_j - \mu_i\|$$

Where μ_i denotes the mean (or centroid) of the cluster i. Initially these means are selected randomly and then at each iteration they are re-computed. K-means provides an efficient way for clustering, even if it suffers from some known drawbacks. The first one is the fact that the number of partitions must be known in advance. Secondly, the idea of generating spherical clusters of similar size is not always appropriate (consider for instance the case of elongated figures or clusters with different size, etc.). For these reasons several alternative versions of the algorithm has been developed (new implementations or integration with other methods), such as the Fuzzy C-means. Ultimately, k-means became a family of algorithms that share the same clustering principle and provides specific adaptations of an elementary technique.

3.5.4 Dominant sets

Dominant sets is a similarity based technique which has elicited an increasing interest in the last few years. It has been built on the assumption that a cluster is defined by two complementary criteria: internal homogeneity (items in the same cluster should be similar to each other) and external in-homogeneity (items in different clusters should be dissimilar). Data to be clustered are represented in an undirected weighted graph $G(V, E, w)$, where $V = \{v_1, v_2, ..., v_n\}$ is the set of nodes, $E \subseteq VxV$ is the set of edges and $w : V \to R_+$ is a positive weight function. In this formal setting, nodes stand for objects and edges express how much object are similar to each other. The main goal of the algorithm is to find a dominant set, that is a coherent set of objects. Intuitively the notion of dominant set can be defined as the largest set of nodes with the highest internal weight. The latter is given by the sum of (relative) pairwise similarities within cluster. In graph theory, this finding corresponds to the notion of a clique, which is a subset of vertices where all nodes are mutually adjacent. In the this context, the idea of dominant set generalizes the concept of a maximal clique, that is a clique in an un-weighted undirected graph, which is not contained in any larger one. Specifically, the technique of dominant sets extends the results of the Motkzin-Strauss theorem, which establishes a correspondence between the maximal/maximum cliques of an un-weighted graph and the local/global solutions of the following quadratic program:

$$\begin{array}{ll} \text{maximize} & f(x) = x^T A x \\ \text{subject to} & e^T x = 1, \quad x \in R^n_+ \end{array}$$

Not that A is the weighted adjacency (or similarity) matrix, where $a_{ij} = w(i, j)$ if $(i, j) \in E$ otherwise $a_{ij} = 0$, and $x$is a n-dimensional vector whose components express the participation of nodes in the cluster: "if a component has a small value, then the corresponding node is weakly associated with the cluster, whereas if it has a large value, the node is strongly associated with the cluster. Components corresponding to nodes not participating in the cluster are zero" [69, p.168]. Hence, roughly speaking, the solution of this quadratic program leads to the discovery of a dominant set. Dominant sets solves this program by using replicator dynamics, a family of continuous and discrete-time dynamical systems arising in evolutionary game theory. In our experiments we used the following discrete version:

$$x_i(t+1) = x_i(t) \frac{(Ax)_i}{x(t)^T A x(t)}$$

Note that algorithm discovers one cluster at time and to find more clusters several solutions are possible. The easiest one is to repeat the algorithm removing objects selected in previous iteration until all objects are assigned. Another one (i.e., the hierarchical version), more successful, is introduced in [76].

3.6 A basic test set

We tested the generalized context model and the rational model on two data sets coming from the repository of the Speech and Image Processing Unit of the school of computing [7]. They consist in a shape data set, which is composed of 373 vectors in 2 dimensions and with 2 clusters and in the well-known Iris data set, which is based on 150 vectors, 4 dimensions and 3 clusters (50 items per groups). Note that both data sets are accompanied by a target partition, that is a classification which provides a term of comparison for algorithm predictions.

For the testing of k-means and dominant sets our starting point was recent work by Pothos and colleagues which sought to compare a number of competing models of spontaneous categorization across a carefully constructed test set of nine different category structures, displayed in Figure [88]. Pothos et al.'s experimental investigation was focused on general properties of categories that might be expected to be relevant to spontaneous categorization, such as the proximity of clusters, number of clusters, relative size of clusters, and cluster spread. The main aim of their investigation was to compare how well different models predicted perceptions of "category intuitiveness", where intuitiveness was operationalized by the extent to which participants agreed on the best partitioning, as measured by the frequency of the preferred classification. Note that, among others, they tested both the generalized context model and the rational model. Moreover the application of the generalized context model, which has been actually thought for supervised categorization, was justified on the theoretical hypotheses that there is no real distinction between supervised and unsupervised classification [87, 2]. Hence, to deal with spontaneous categorization Pothos and colleagues proposed an unsupervised version of that model.

3.7 Some experimental results

In our experiments we used a classification error function measuring the number of items incorrectly classified with respect to the target partition. With respect to the stimulus sets of cognitive psychology we compared the behaviour of pattern recognition algorithms against the results reported in [88] , repeating the test 20 times. So, the accuracy reported is based on the mean of all these runs.

The results on the test set designed by Pothos et al. are reproduced in the figure 3.3. Experiment results are obviously influenced by algorithm settings. Indeed they do start from very different points (for instance, in k-means the number of classes is a priori defined and in GCM we have one labelled item per group, etc.). But, despite these differences we treated them as plausible models of spontaneous categorization.

What emerges from such tests is that none of the proposed methods is able to predict correctly the ordinal properties of nine different category structures. Indeed

[7] This is the web site: `http://cs.joensuu.fi/sipu/datasets/`

Figure 3.2: Pothos et al.s stimulus sets along with the preferred partitions. Reproduced from [88] with permission.

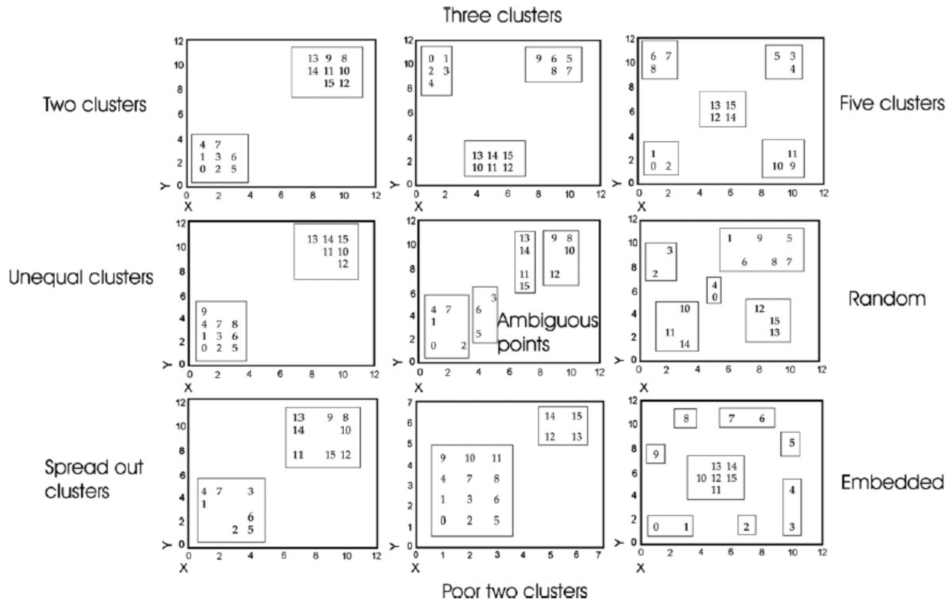

none really reproduce the same behaviour of human preferences, traced by the black line. According to the empirical observations the "five-cluster stimulus set" is the best predicted and such preference is well distinguished from the rest of stimulus sets[8]. While dominant sets an the generalized context model, for instance, predict well the "five-cluster stimulus set", the "two-cluster stimulus set" and others with no distinction.

To test the models of cognitive psychology on the shape and the Iris data sets we used the same error measure. The numerical results are referred to the mean error produced across 3 algorithms simulations. With rational model we set the coupling parameter with the value which produced the best fit in the experimental investigation presented in [88], i.e. $c = 1/3$. In our experiments we tested both the original Anderson's algorithm (Local MAP) and the more recent approximation by Gibbs sampling presented in [98]. While for testing the generalized context model on the pattern recognition data sets we employed a semi-supervised version since the (unsupervised) implementation proposed in [88] had elevated computational cost in dealing with hundreds of items. Then, the parameters were fixed as follows: the weights of feature dimensions, $w_i = 1/d$, with $d =$ number of dimensions; the form of distance metric $r = 2$ (Euclidean distance); the shape of the function $f = 2$ (gaussian); and sensitivity parameter $c = 22$ (for Iris data set) , 10 (for Jain data set).

[8] Note, indeed, that the stimulus sets are ordered on the basis of the perceived intuitiveness.

Figure 3.3: Behaviour of algorithms with respect to the human response on the nine stimulus sets used in [88].

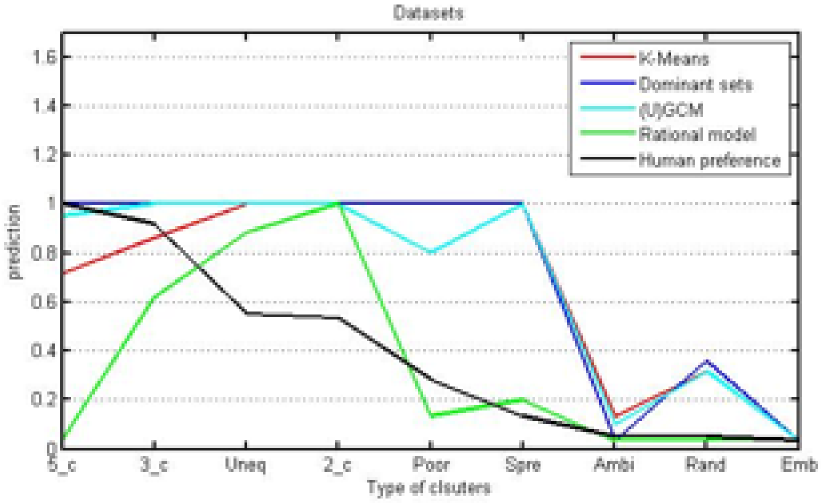

Table 3.1: Results with respect to pattern recognition data sets

Class. Error	dominant sets	k-means	GCM	rational model
Iris	0.0333	0.1743	0.1400	0.333(Local Map),0.2669(Gibbs)
Jain	0.3834	0.2172	0.0342	0.261(Local Map),0.4906(Gibbs)

Models' performance are reported in the table 3.1 [9]. In some circumstances the generalized context model predicted correctly all assignments (see, for instance, the plot of semi-supervised GCM on the shape data set in figure 3.4). Note that this result was not so influenced by the values of free parameters, but rather the choice of labelled items. This could suggest that the optimization process used to calculate the value of parameters [88] could be avoided if we found good exemplars, in the sense that they better propagate similarity information. Also in the case of Iris data set predictions are quite good, producing a mean error less than that produced by k-means.

With respect to the rational model we found significant differences in predictions produced by the local MAP and the Gibbs sampling algorithm. Specifically, if we look at the Iris data set we will find that the Local MAP algorithm identifies correctly the linearly separable cluster (Iris Setosa) and collects the other two groups (Iris Versicolor and Iris Virginica) in one single class. On the other hand, Gibbs sampling seems to provide a more articulated partition, where the rest of non linearly separable classes (Iris Versicolor and Iris Virginica) is split in three clusters (see figure 3.5). We observed a similar behaviour on the Jain data set as well. In this case, Local Map included all objects in one group, while Gibbs sampling tended to divide the

[9] Note that with respect to Iris data set the best fit was obtained by using an hierarchical implementation of dominant sets citePavPel03

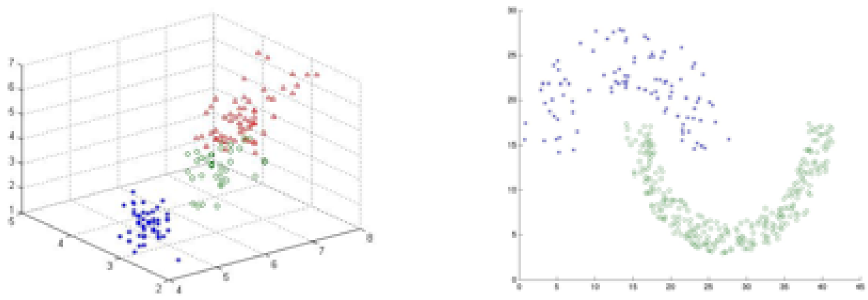

Figure 3.4: GCM on Iris (on the left) and Jain (on the right) data sets

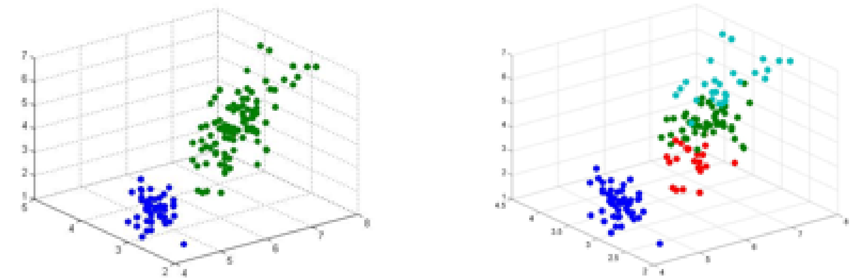

Figure 3.5: Figure n. 7: Local Map (on the left) and Gibbs Sampling (on the right)

two main groups in smaller parts (see figure 3.6).

3.8 Discussion

During our excursus we encountered two distinct perspectives on categorization. Both of them provide a certain view on the problem and somehow answer some fundamental questions such as those concerning the nature of objects to be categorized, the structure of the process of categorization and the global characterization of the issue. In section 3.4 we collected some partial statements around these points of interactions by reasoning on descriptions of how pattern recognition and cognitive psychology approach the theme. Now, we would like to extend the initial evaluation even through the lens of our practical comparison.

- *Object of categorization.* As we saw before, in cognitive psychology categorization is directed towards "plain" objects (with respect to the structure and to the number of instances), the majority of which is built for experimental purposes. Cognitive psychology, indeed, designs data sets and experiments primarily to discover something about mind not about objects. This intention leads the field to constrain the space of possible solutions not by introducing

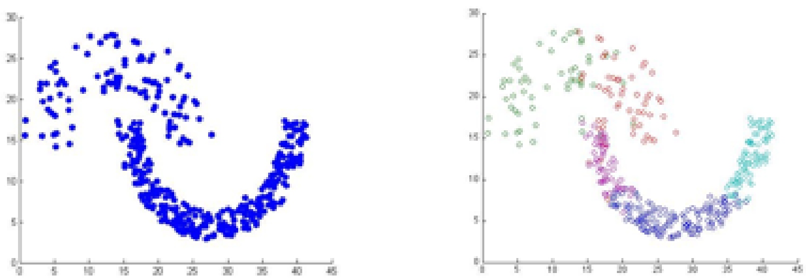

Figure 3.6: Figure n. 8: Local Map (on the left) and Gibbs Sampling (on the right)

special indications to the classifier (either human or automated) but by reducing the complexity of data. In so doing, psychology can isolate the points of interest for the inquiry (e.g.: proximity among clusters, tightness of clusters, tec.) and leave aside those aspects which may affect the complexity of clustering choices (e.g.: aspects involving multiple personal or cultural biases). However, the employment of some pattern recognition data sets might have several positive effects on cognitive psychology. First, it could enhance the critical evaluation of models increasing the opportunities for controlling theories or experimental activities [10], in the sense implied by Poppers falsificationism [83]. This may hold, for instance, when psychologists put forward some ideas which could make sense to some extent, but under another perspectives might lose their feasibility or at least part of their strength. An example has been shown in the section of experiments, where the unsupervised version of the generalized context model has been revised in order to deal with pattern recognition data sets. In that case we could not apply the method proposed in [88] directly, since it was computationally expensive (to reach the target partition, in fact, it may be forced to explore all possible assignments). Therefore the supposed equivalence of unsupervised and supervised methods was in a sense "falsified" by our tests. Indeed we realized that a direct employment of the generalized context model in unsupervised classification was not feasible on large data sets and further alternatives should be considered. For instance, in our case we chose the option of a semi-supervised implementation. Another benefit that can be derived from interactions with machine learning is the introduction of computational constraints. Bayesian approaches, like rational model, are focused on defining the optimality of a cognitive process without much interest in computational costs (in terms of running time). Recently, Markman and Otto [63] noticed that the key limitation of "Bayesian

[10]In this respect for cognitive psychology could be interesting to investigate different level of abstraction of objects. And, for instance, a possible question might be: to what extent an experiment design should abstract object so as to not lose the applicability to some realistic categorization scenario?

fundamentalism" is that "it focuses selectively on optimality of information processing rather than on the combination of information and time" [63, p. 207]. Sanborn, Navarro and Griffiths proposed a general strategy to overcome the cost of Bayesian inference. However the approximation they proposed by means of Gibbs sampling is time-consuming over Iris and Jain data sets. In fact the process to generate samples from the desired probability distribution becomes more and more onerous as the number of objects increases. Moreover in the computation one should consider some waste samples, because of the burn-in iterations (early samples not yet coming from the desired distribution) and the space-sampling (due to the strong dependency that might be between one iteration to the next one).

- *Process of categorization.* In the description of scientific practices we discovered a fundamental distinction between cognitive psychology and pattern recognition. From the side of cognitive psychology we draw the idea of a process strictly related to the agent, while on the side of pattern recognition the reference to the "actor" of categorization is much more shaded. This has several consequences. The first one regards the possibility to create connections with a wide psychological and philosophical literature, which has a long tradition in the theme. Cognitive psychology is much more used to take cues from such tradition and most of its models might be seen as a technical extension of what has been explored in other ways. In this context formalisms and algorithms make sense just in relation to some theoretical hypothesis, and between these two levels (practical - theoretical) there is a continuous exchange of information. Differently, pattern recognition does not necessarily need to bring its methods back to a further level of explanation. It rather asks for reliable solutions to practical problems of pattern discovery (e.g.: spam filtering, face detection, topic spotting, customer segmentation, etc.) in greater accordance with an engineering perspective. Another effect is visible in the possible approach to validation. In the field of pattern recognition there are multiple options. Sometimes the ground truth is defined by expert people or derived by some specific knowledge (e.g.: biology or physics). But there are other situations in which it could make sense understanding algorithm output in the light of psychological experiments. For example, when we test algorithms on abstract data or without particular expectations (e.g.: we do not know the application domain and we charge the algorithm to find some patterns or regularities) it may be interesting to compare predictions with human behaviour. In some cases this happens but such experiences are mostly isolated initiatives. In this way one could integrate independent measure of goodness (e.g. inner quality score or "object-oriented" approaches) with human-dependent preferences, that is, with evidence gathered during specific lab tests. We know that clustering validation is a relevant debate in the community of machine learning. A recent critique supported the idea of an application-dependent evaluation (see [118]) since it seems impossible to abstract clustering problems from the end-user intent. To this end a "taxonomy of clustering problems to identify clustering applications" has been proposed. We think that in this taxonomy there might be space for psychological intents as well considering the fact that, like many engineering projects, they arise out of some specific problems whose complexity lies in a different level of complexity. Indeed, as we saw so far, cognitive psychology encounters complexity not really in solving a specific clustering task (e.g.: grouping e-mails or costumers, etc.) but, rather, in the modeling a clustering scenario (i.e., a more abstract and general experience).

- *Theory of categorization.* Let us to come back to some initial points: are cognitive psychology and pattern recognition modelling just the same phenomenon? or are they addressing different topics? A shallow judgment may consider these two areas as facing the same space of problems, as they were almost exchangeable. But a deeper insight shows that today's research in pattern recognition differ from cognitive psychology in several aspects. We have already suggested that pattern recognition covers both scientific and engineering issues. This point has been stressed also by some recent reflections on this research areas. And for instance, Duin and Pekalska acknowledged that "automatic pattern recognition is usually considered as an engineering area which focuses on the development and evaluation of systems that imitate or assist humans in their ability of recognizing patterns. It may, however, also be considered as a science that studies the faculty of human beings (and possibly other biological systems) to discover, distinguish, characterize patterns in their environment and accordingly identify new observations" [31, p.]. However it is undeniable that, compared to cognitive psychology, the mainstream of today's pattern recognition gives a different emphasis to the engineering and the scientific aspects. Indeed it seems that the engineering factors are much more relevant in pattern recognition whereas the attempts to "imitate" human behaviour clearly prevail in the field of cognitive psychology. But as we discussed in the chapter 1 the engineering activity is often continuous with ordinary science so that a special engineering attitude in pattern recognition should not necessarily mean a lack of scientific work. So, in what sense does the scientific dimension of pattern recognition differ from the scientific activity of cognitive psychology? Maybe the difference may be referred to the way in which the relationship between a problem and its solution is addressed. In pattern recognition, for instance, what seems to be less considered is the fact that the reason of solving in science is ultimately the search for knowledge and explanations. Indeed even Laudan who advocated that science is a problem-solving activity admitted that such a type of activity starts from "anything about the natural world which strikes us as odd, or otherwise in need of explanation" [60, p. 15]. This could mean that the success or the failure of a technique are included in a wider research context which goes beyond the particular application. For instance, with respect to our experiments, for instance, we could be interested to understand why two approximations (i.e., Local Map and Gibbs Sampling) which ought to be in principle equivalent express such a divergence or what is the (theoretical) meaning of a technical adaptation for a specific method (e.g., the hierarchical version of dominant set with respect to the original method).

3.9 Summary

We introduced how the problem of categorization is addressed by two similar but distinct research areas, that is cognitive psychology and pattern recognition. We pointed out their specific approaches focusing three main points: the object of categorization, the process of categorization and the theory of categorization. Our investigation showed that there exist several differences both in the theoretical (e.g.: how they design experiment, how they conceive the data to be clustered, or what is the meaning of evaluation phase) and in the experimental evaluations (i.e. the performance of methods). On the one hand, the field of pattern recognition seem to prevail on the experimental test because no strong adjustments were required to make its methods suitable for cognitive data sets. On the other hand, cognitive psychology benefits from its ability to build model and to formulate explanation of what is predicted allowing researchers to arrange the inquiry from the very beginning (i.e. exploring issues which go beyond the technical results). Hence, in consideration of some theoretical and experimental observations we outlined the two fields interact. Future work might be directed toward a deeper investigation of such interdependencies, focusing specific aspects, such as the design of experiments, or specific models, such as the psychological plausibility of specific pattern recognition methods.

4

How mature is the field of pattern recognition?

4.1 The theoretical perspective

According to Thomas Kuhn, the "acquisition of a paradigm and of the more esoteric type of research it permits is a sign of maturity in the development of any given scientific field" [53, p. 11]. In this paper, we propose to address the question whether the fields of pattern recognition and machine learning have achieved the level of maturity in the sense implied by the quotation above.[1] Note that Kuhn's notion is quite different from (and indeed more profound than) the commonsensical view which maintains that "mature scientific disciplines are expected to develop experimental methodologies, comparative evaluation techniques, and theory that is based on realistic assumptions" [47, p. 112]. Under this interpretation, one would be tempted to respond with an emphatic "yes" to the question posed above. Indeed, in the last 25 years or so researchers have dramatically changed their attitude to the evaluation of new algorithms and techniques, and it seems that Langley's well-known incitement to make machine learning an "experimental science" [57] has been taken seriously by the community. Since its birth in the late 1990's, for example, the UCI ML repository keeps growing at a fast pace and at the time of writing it contains 244 different data sets on the most disparate problems and applications. On the other hand, there is an increasing level of sophistication in the way in which the performance of the algorithms are quantitatively evaluated, and we saw an evolution from simple scalar performance measures such as the classification accuracy to more elaborated ones such as ROC curves and statistical tests. However, a deeper analysis reveals that the situation is more controversial than it appears, as there is more to science than simply experimental analysis, and the equation "scientific = experimental" is too naive to satisfactorily capture the multifaceted nature of the "scientific method" (granted that there exists one [33]).

In this paper, however, we do not intend to enter into this discussion, but would

[1] A note on terminology: although throughout the paper we consistently use the term "pattern recognition," we in fact think that much of the discussion could be referred also to the field of machine learning.

like to attack the question from a purely Kuhnian perspective according to which, as anticipated above, the notion of maturity in science is inextricably tied to the concept of a "paradigm," one of the cornerstones of twentieth-century philosophy of science. In one sense, our motivating question can in fact be interpreted simply as an excuse to analyze the current status of the pattern recognition and machine learning fields using the conceptual tools provided by Kuhn. Note that under this interpretation, there is no pretense of judging the "scientificity" of a given research area or to provide a demarcation, à la Popper, between scientific and non-scientific fields. Indeed, using Kuhn's suggestion, Aristotelian physics, which dominated the scene for over two millennia, has to be considered as mature as today's physics although of course, according to the modern interpretation of the term, we would not dream of calling it a "science." [2]

The publication of Kuhn's *Structure of Scientific Revolutions* in 1962 was a momentous event in the modern history of ideas. It provoked itself a revolution in the way we think at science whose far-reaching effects are felt in virtually all academic as well as popular circles. What made Kuhn's image particularly successful in describing the nature of scientific progress is, no doubt, his notion of a paradigm. Unfortunately, the reception of the term by the philosophical and scientific communities was controversial, and a number of difficulties persuaded Kuhn to clarify his position in his famous 1969 *Postscript* [53]. Indeed, as he himself admitted, the term was used in a vague and ambiguous way throughout the book but, besides minor stylistic variations, he identified two very different usages of the term. On the one hand, he aimed to describe some accepted examples which serve as a model for the solution of new puzzles (the "narrow" sense), whereas, on the other hand, he meant a more profound commitment to a set of beliefs and values (the "broad" sense).

In this chapter, we aim to approach the question posed in the title by exploiting both interpretations of the concept, and the discussion will make it clear that the answer depends on which sense one considers. Note that Cristianini [22] has recently undertaken a study similar in spirit to ours, but he seems to have emphasized mostly the first, narrow, sense of the term. In contrast, we shall focus more on the broad interpretation. This will give us the opportunity to discuss the philosophical (often tacit) assumptions underlying much of contemporary pattern recognition research and to undertake a critical reflection of its current status. In particular, we will see how deep is the bond with essentialism, one of the oldest and most powerful ideas in the whole history of philosophy, and we shall maintain that the community is gently moving away from it, a phenomenon which, we shall speculate, seems to parallel the rejection of the essentialist hypothesis by modern science.

[2] See Kuhn's autobiographical fragment contained in [54] for a rehabilitation of Aristotle as a physicist.

4.2 Kuhnian analysis in pattern recognition

Before undertaking our exploration we need to motivate the pertinence of Kuhn's contribution within the context of pattern recognition This is required by the fact that Kuhn's investigation is essentially directed towards well-established sciences such as physics or chemistry. Thus, it would seem sensible to ask whether pattern recognition is an appropriate subject of study under a Kuhnian approach.

We feel the question of particular relevance in view of the fact that the context in which pattern recognition research has grown up is interdisciplinary in a profound way. Such a distinct character of the evolution of the field has not always been recognized by researchers so that the aspects involved are usually seen in isolation. This resulted in the common reductionist tendency to conceive the fields of pattern recognition machine learning as either engineering or science. Some scholars, indeed, assume that, as well as providing technical solutions, the fields of pattern recognition and machine learning deal with fundamental questions pertaining to categorization, abstraction, generalization, induction, etc., and, in so doing, their contribution is in fact scientific [31, 120]. In some cases, the approach of machine learning has been associated even to the scientific practice of physics [105] or, more generally, to experimental sciences [57]. Conversely, nowadays it prevails the idea that these areas are primarily engineering disciplines. For example, Pavlidis, one of the pioneers of the field, recalling Hermann Hesse's novel *Das Glassperlenspiel* claims that "the prospects are bright if we approach pattern recognition as an engineering problem and try to solve important special cases while staying away from the Glassperlenspiel. The prospects are grim if we keep looking for silver bullets that will solve "wholesale" a large range of general problems, especially if we harbor the illusion of doing things the way the human brain does."[77, p. 7]. And this idea has been more recently echoed by von Luxburg et al. [118].

This sharp opposition between science and technology stems from an oversimplified view of their mutual relationship. However, in the light of some new achievements in the philosophy of technology (see, e.g., [34]), it turns out that, granted that there are indeed important differences, at the conceptual level the boundary between the two camps is more blurred than is commonly thought, and that they stand to each other in a kind of circular, symbiotic relationship. Indeed, technology can be considered as an activity producing new knowledge on a par with ordinary science. The so called operative theories [16] in technology look like those of science and their contribution goes beyond the mere application of scientific knowledge. On the other hand, even science can be brought closer to technology when its progress is expressed in terms of "immanent goals." This idea lies at the heart of Laudan's problem-solving approach to science [59] and could well characterize much of the work in the field of pattern recognition.

The profound interaction between scientific and technological components is a key to understand the pattern recognition activity and other research areas within artificial intelligence. The history of the field, in fact, counts numerous examples of this fecund relationship. The case of neural networks is particularly significant, as their original formulation had a genuine scientific motivation, that is, the wish of studying and imitating the brain but, in the phase of their renaissance, technical matters prevailed. Indeed, with the (re)invention of the back-propagation algorithm for multi-layer neural networks and, above all, thanks to the impressive results obtained by these new models on practical problems such as zip code recognition and speech synthesis, a new wave of excitement spread across the artificial intelligence community. At that point, however, it was already clear that these models had no pretense of being biologically plausible [21]. Pavlidis nicely summed up this state of affairs by noting that "the neural networks that have been in vogue during the last 15 years may be interesting computational devices but they are not models of the brain. (Except maybe of the brains of people who make that claim sincerely)" [77, p. 2]. Bayesianism is another interesting example of the gate allowing pattern recognition to move from theoretical issues to more practical aims. Introduced as a theory which can characterize the strength of an agent's belief, it provided many inference algorithms with a practical machinery. On the other hand, recent advances in density estimation techniques, such as nonparametric Bayesian methods, have been successfully applied to approach a variety of cognitive processes [98].

To sum up, the recent contributions of philosophy of technology and of the philosophy of science lead us to rethink the classical dichotomy between science and technology, which is still holding in some subfields of artificial intelligence, as they appear closer than we used to think. Historical examples suggest that machine learning and pattern recognition work, indeed, as a bridge between the two and many ideas from science result in technological innovation and vice versa [81]. In reference to our discussion, this means that a contribution from philosophy of science should not be considered irrelevant for these two fields since the scientific side is as much important as the technological one. Accordingly, we do think that Kuhn's analysis is not only appropriate to the pattern recognition research but could also contribute to get a deeper understanding of its nature.

4.3 Kuhn's notion of a paradigm

In his *Structure* (as the book is known) [53], Kuhn provides an account of scientific development that is dramatically different from the standard idea of a steady, cumulative progress. According to him, a science traverses several discontinuities alternating "normal" and "revolutionary" phases. During normal periods the development of a science is driven by adherence to a "paradigm" whose function is to support scientists in their "puzzle-solving" activity with a number of practical and theoretical tools, including theories, values and metaphysical assumptions. When some worrying puzzles remain unsolved (the so-called "anomalies") and the current approach loses progressively its original appeal, a discipline enters a period of crisis. At this point, the activity is characterized by "a proliferation of competing articulations, the willingness to try anything, the expression of explicit discontent, the recourse to philosophy and to debate over fundamentals"[53, p. 91]. Finally, the crisis is resolved by a scientific revolution leading to the replacement of the current paradigm by a new one. The revolution results in a paradigm shift, after which a discipline returns to a normal phase,based this time on a new accepted framework.

As we have seen before, there are two distinct uses of the notion of a paradigm. At first, Kuhn uses the term "paradigm" to refer to some concrete achievements that can work as models or examples and supply explicit rules for the solution of the remaining puzzles. In the history of science examples of this notion abound and include, e.g., Newton's mechanics and Franklin's theory of electricity, which implicitly defined the legitimate problems and methods of a research field for succeeding generations of practitioners. A second way to apply the term "paradigm" refers to a more global sense and includes, above all, concepts, theoretical principles, metaphysical assumptions, worldviews, etc.

In his *Postscript*, Kuhn introduced the idea of a broad paradigm in terms of a "disciplinary matrix", which could be seen as a theoretical and methodological framework wherein scientists conduct their research. This framework includes the basic assumptions of a discipline providing a community with the practical and theoretical indications, for instance, about how to lead investigations or what to expect from experiments. Among the elements which compose this matrix, we aim to focus on symbolic generalizations and metaphysical paradigms.

A research community could easily present formal expressions or codified terms that could live in the acceptance of all members for several years. These are what Kuhn calls "symbolic generalizations" and their function goes basically in two directions. That is, they can work as laws of nature, such as $\mathbf{F} = m\mathbf{a}$, or "elements combine in constant proportion by weight" [53, p. 183], but they can also serve to settle some fundamental definitions assigning symbols to specific meanings. Note that, according to Kuhn "all revolutions involve, among other things, the abandonment of generalizations the force of which had previously been in some part that tautologies"[53, p. 184].

A second type of component is given by the metaphysical parts of a paradigm. Metaphysical elements can be beliefs or models and incorporate tacit or implicit knowledge. In practice these components shape the general disposition and the methodological attitude of a scientist suggesting particular metaphor or world-views. The strength of such components is that of determining what will be accepted as as an explanation and, above all, the importance of unsolved puzzles.

4.4 Paradigms in pattern recognition: The broad perspective

Forms of narrow paradigms can be easily found in pattern recognition research. The evolution of the field, indeed, is a story of great achievements that were able to create strong traditions around them. An obvious example is provided by neural networks which played a key role in the early as well as later developments of the field. More recent examples include, e.g., kernel methods and spectral clustering. Some of these success stories are collected in [22], which nicely describes the transition from the "knowledge-driven" to the"learning-driven" paradigm in artificial intelligence.

4.4.1 The disciplinary matrix

Here we would like instead to focus on the broad sense of the notion of a paradigm and see whether the contours of a disciplinary matrix come up through the concrete practice of the discipline within the community. Hence, with Kuhn, we could ask: "what do its members share that accounts for the relative fulness of their professional communication and the relative unanimity of their professional judgments?"[53, p. 182]. To address this issue we will consider the components presented above: the symbolic generalization and the metaphysical paradigm.

First, note that the majority of traditional pattern recognition techniques are centered around the notion of "feature"[25, 12]. Indeed, within the field there is a widespread tendency to describe objects in terms of numerical attributes and to map them into a Euclidean (geometric) vector space so that the distances between the points reflect the observed (dis)similarities between the respective objects. This kind of representation is attractive because geometric spaces offer powerful analytical as well as computational tools that are simply not available in other representations. In fact, classical pattern recognition methods are tightly related to geometrical concepts and numerous powerful tools have been developed during the last few decades, starting from linear discriminant analysis in the 1920's, to perceptrons in the 1960's, to kernel machines in the 1990's.

In the light of Kuhn's perspective we could think of such a representational attitude in terms of a collection of symbolic generalizations which lead the community to take some definitions or principles for granted. Indeed, the development of the field has been accompanied by the deployment of codified terms such as "feature extraction," "feature vector," "feature space," etc., and even by a formal vocabulary which is the basis of the subsequent mathematical manipulation. As a whole, symbolic generalizations have contributed to the general acceptance of a clear idea of what categories are and how they do form, that is the conviction that a classifier groups a set of objects under the same label because of some common features.

But the content of such generalizations might be read also at the level of the metaphysical paradigm. This brings us to discuss the philosophical assumptions behind pattern recognition and machine learning research. In fact, as pointed out in [25], their very foundations can be traced back to Aristotle and his mentor Plato who were among the firsts to distinguish between an "essential property" from an "accidental property" of an object, so that the whole field can naturally be cast as the problem of finding such essential properties of a category. As Watanabe put it [120, p. 21]: "whether we like it or not, under all works of pattern recognition lies tacitly the Aristotelian view that the world consists of a discrete number of self-identical objects provided with, other than fleeting accidental properties, a number of fixed or very slowly changing attributes. Some of these attributes, which may be called "features", determine the class to which the object belongs." Accordingly, the goal of a pattern recognition algorithm is to discern the essences of a category, or to "carve the nature at its joints." In philosophy, this view takes the name of *essentialism* and has contributed to shape the puzzle-solving activity of pattern recognition research in such a way that it seems legitimate to speak about an essentialist paradigm.

4.4.2 Essentialism and its discontents

Essentialism has profoundly influenced most of scientific practice until the nineteenth century, even though early criticisms came earlier with the dawn of modern science and the new Galilean approach. Later, William James, deeply influenced by Darwin, went so far as to argue that "[t]here is no property ABSOLUTELY essential to any one thing. The same property which figures as the essence of a thing on one occasion becomes a very inessential feature upon another" [48, p. 959]. Nowadays, anti-essentialist positions are associated with various philosophical movements including pragmatism, existentialism, decostructionism, etc., and is also maintained in mathematics by the adherents of the structuralist movement, a view which goes back to Dedekind, Hilbert and Poincaré, whose basic tenet is that "in mathematics the primary subject-matter is not the individual mathematical objects but rather the structures in which they are arranged" [92, p. 201]. Basically, for an anti-essentialist what really matters is relations, not essences. The influential American philosopher Richard Rorty nicely sums up this "panrelationalist" view with the suggestion that there are "relations all the way down, all the way up, and all the way out in every direction: you never reach something which is not just one more nexus of relations." [93, p. 54]

During the 19th and the 20th centuries, the essentialist position was also subject to a massive assault from several quarters outside philosophy, and it became increasingly regarded as an impediment to scientific progress. Strikingly enough, this conclusion was arrived at independently in at least three different disciplines, namely physics, biology, and psychology.

In physics, anti-essentialist positions were held (among others) by Mach, Duhem, Poincaré, and in the late 1920's Bridgman, influenced by Einstein's achievements, put forcefully forward the notion of operational definitions precisely to avoid the troubles associated with attempting to define things in terms of some intrinsic essence [14]. For example, the (special) theory of relativity can be viewed as the introduction of operational definitions for simultaneity of events and of distance, and in quantum mechanics the notion of operational definitions is closely related to the idea of observables. This point was vigorously defended by Popper [85], who developed his own form of anti-essentialism and argued that modern science (and, in particular, physics) was able to make real progress only when it abandoned altogether the pretension of making essentialist assertions, and turned away from "what-is" questions of Aristotelian-scholastic flavour.

In biology, the publication of Darwin's Origin of Species in 1859 had a devastating effect on the then dominating paradigm based on the static, Aristotelian view of species, and shattered two thousand years of research which culminated in the monumental Linnaean system of taxonomic classification. According to Mayr, essentialism "dominated the thinking of the western world to a degree that is still not yet fully appreciated by the historians of ideas. [...] It took more than two thousand years for biology, under the influence of Darwin, to escape the paralyzing grip of essentialism."[65, p.87]

More recently, motivated by totally different considerations, cognitive scientists have come to a similar discontent towards essentialist explanations. Indeed, since Wittgenstein's well-known family resemblance argument, it has become increasingly clear that the classical essentialist, feature-based approach to categorization is too restrictive to be able to characterize the intricacies and the multifaceted nature of real-world categories. This culminated in the 1970's in Rosch's now classical "protoype theory" which is generally recognized as having revolutionized the study of categorization within experimental psychology; see [56] for an extensive account, and [118] for a recent evocation in the pattern recognition literature.

The above discussion seems to support Popper's claim that every scientific discipline "as long as it used the Aristotelian method of definition, has remained arrested in a state of empty verbiage and barren scholasticism, and that the degree to which the various sciences have been able to make any progress depended on the degree to which they have been able to get rid of this essentialist method" [84, p. 206].

4.5 Signs of a transition?

It is now natural to ask: what is the current state of affairs in pattern recognition? As mentioned above, the field has been dominated since its inception by the notion of "essential" properties (i.e., features) and traces of essentialism can also be found, to varying degrees, in modern approaches which try to avoid the direct use of features (e.g., kernel methods). This essentialist attitude has had two major consequences which greatly contributed to shape the field in the past few decades. On the one hand, it has led the community to focus mainly on feature-vector representations. On the other hand, it has led researchers to maintain a reductionist position, whereby objects are seen in isolation and which therefore tends to overlook the role of relational, or contextual, information.

However, despite the power of vector-based representations, there are numerous application domains where either it is not possible to find satisfactory features or they are inefficient for learning purposes. This modeling difficulty typically occurs in cases when experts cannot define features in a straightforward way (e.g., protein descriptors vs. alignments), when data are high dimensional (e.g., images), when features consist of both numerical and categorical variables (e.g., person data, like weight, sex, eye color, etc.), and in the presence of missing or inhomogeneous data. But, probably, this situation arises most commonly when objects are described in terms of structural properties, such as parts and relations between parts, as is the case in shape recognition [11]. This led in 1960's to the development of the structural pattern recognition approach, which uses symbolic data structures, such as strings, trees, and graphs for the representation of individual patterns, thereby, reformulating the recognition problem as a pattern-matching problem.

Note that, from a technical standpoint, by departing from vector-space representations one is confronted with the challenging problem of dealing with (dis)similarities that do not necessarily possess the Euclidean behavior if there exists a configuration of points in some Euclidean space whose interpoint distances are given by D, or not even obey the requirements of a metric. The lack of the Euclidean and/or metric properties undermines the very foundations of traditional pattern recognition theories and algorithms, and poses totally new theoretical/computational questions and challenges. In fact, this situation arises frequently in practice. For example, non-Euclidean or non-metric (dis)similarity measures are naturally derived when images, shapes or sequences are aligned in a template matching process. In computer vision, non-metric measures are preferred in the presence of partially occluded objects [45]. As argued in [45], the violation of the triangle inequality is often not an artifact of poor choice of features or algorithms, and it is inherent in the problem of robust matching when different parts of objects (shapes) are matched to different images. The same argument may hold for any type of local alignments. Corrections or simplifications may therefore destroy essential information.

As for the reductionist position, in retrospect it is surprising that little attention has typically been devoted to contextual information. Indeed, it is a common-sense observation that in the real world objects do not live in a vacuum, and the importance of context in our everyday judgments and actions can hardly be exaggerated, some having gone so far as to maintain that all attributions of knowledge are indeed context-sensitive, a view commonly known as contextualism [90]. Admittedly, the use of contextual constraints in pattern recognition dates back to the early days of the field, especially in connection to optical character recognition problems and it reached its climax within the computer vision community in the 1980's with the development of relaxation labeling processes and Markov random fields [44]. However, all these efforts have soon fallen into oblivion, mainly due to the tremendous development of statistical learning theory, which proved to be so elegant and powerful. Recently, the computer vision community is paying again increasing attention to the role played by contextual information in visual perception, especially in high-level problems such as object recognition (see, e.g., [74]), and neuroscientists have started understanding how contextual processing takes actually place in the visual cortex.

It is clearly open to discussion to what extent the lesson learnt from the historical development of other disciplines applies to machine learning and pattern recognition, but it looks at least like that today's research in these fields is showing an increasing propensity towards anti-essentialist/relational approaches (see [78, 79] for recent accounts). Indeed, in the last few years, interest around purely similarity-based techniques has grown considerably. For example, within the supervised learning paradigm (where expert-labeled training data is assumed to be available) the now famous "kernel trick" shifts the focus from the choice of an appropriate set of features to the choice of a suitable kernel, which is related to object similarities [106]. However, this shift of focus is only partial as the classical interpretation of the notion of a kernel is that it provides an implicit transformation of the feature space rather than a purely similarity-based representation. Similarly, in the unsupervised domain, there has been an increasing interest around pairwise algorithms, such as spectral and graph-theoretic clustering methods, which avoid the use of features altogether [107, 69]. Other attempts include Balcan et al.'s theory of learning with similarity functions [10], and the so-called collective classification approaches, which are reminiscent of relaxation labeling and similar ideas developed in computer vision back in the 1980's (see, e.g., [104] and references therein).

Despite its potential, presently the similarity-based approach is far from seriously challenging the traditional paradigm. This is due mainly to the sporadicity and heterogeneity of the techniques proposed so far and the lack of a unifying perspective. On the other hand, classical approaches are inherently unable to deal satisfactorily with the complexity and richness arising in many real-world situations. This state of affairs hinders the application of pattern recognition techniques to a whole variety of real-world problems. Hence, progress in similarity-based approaches will surely be beneficial for pattern recognition as a whole and, consequently, for the long-term enterprise of building "intelligent" machines.

4.6 Summary

How are we to respond to the question which motivated the present study? Clearly the answer depends on the scope of the notion of a paradigm chosen (narrow *vs.* broad). If we stick to the narrow interpretation, we easily arrive at the conclusion, with Cristianini [22], that the fields of machine learning and pattern recognition are indeed mature ones, so much so that in their (short) history we have had a whole succession of paradigms, intended as specific achievements which attracted the attention of a large and enduring fraction of the community.

Our study, however, focused on the broad interpretation of the term and this led us to discuss the philosophical underpinnings of much of contemporary pattern recognition research. Our analysis has shown that the community has traditionally adhered, by and large, to an essentialist worldview, where objects are characterized and represented in terms of intrinsic, essential features. This view has long been abandoned by modern science and has been in fact considered an impediment to its development. *Mutatis mutandis*, nowadays we are witnessing an increasing discontent towards essentialist representations in the pattern recognition community [78, 79]. Hence, although using Kuhn's view, we might say that the field has reached a satisfactory level of maturity even using his broad interpretation, there are signs which make us think that there is a need to bring to full maturation a paradigm shift that is just emerging, where researchers are becoming increasingly aware of the importance of similarity and relational information *per se*, as opposed to the classical feature-based (or vectorial) approach. Indeed, the notion of similarity (which appears under different names such as proximity, resemblance, and psychological distance) has long been recognized to lie at the very heart of human cognitive processes and can be considered as a connection between perception and higher-level knowledge, a crucial factor in the process of human recognition and categorization [38].

We conclude by noticing that according to Kuhn's picture, at any particular time a scientific field is supposed to have only *one* paradigm guiding it. Applied to pattern recognition, this interpretation seems too restrictive as none of the paradigms mentioned above (either broad or narrow) has really guided the research of the *whole* community. More recent developments of Kuhn's thought, which allow for multiple competing paradigms per time, can be found in Lakatos' and Laudan's work who talked about "research programmes" or "research traditions," respectively [55, 59]. It is therefore tempting to explore whether, in order to provide a more faithful picture of the status of the pattern recognition field, we need to resort to these more sophisticated conceptual tools.

Conclusions

In this thesis we have tried to investigate some philosophical issues arising in pattern recognition research. The philosophical work has been done basically by a critical examinations of some aspects of pattern recognition activity. First we addressed the problem of the nature of pattern recognition pointing out in which sense this research area could be considered as both science and engineering. This has been done thanks to the recent contributions of the philosophy of science and the philosophy of technology which provided our discussion with specific analytical tools (respectively the notions of problem-solving approach and technological theories). The second step was devoted to the traditional formulation of a pattern recognition problem. Symbols and terminology were highlighted by the philosophical analysis of induction which lies, indeed, at the very heart of pattern recognition problems. Then some relationship between pattern recognition and other neighbouring areas were examined (i.e. artificial intelligence and machine learning). Indeed it is not always so clear how pattern recognition is related to artificial intelligence (is pattern recognition a sub-field of artificial intelligence?) and to machine learning (do they are the same?). In the chapter 3 we moved towards a concrete comparison between pattern recognition and cognitive psychology. Indeed both pattern recognition and cognitive psychology may be considered thorough the lens of categorization. But, even though pattern recognition and cognitive psychology share common problems (and often also common language and formalism) they carry on two different problem-solving activities. At first sight, this might seem actually obvious but a deeper evaluation, by contrast, has revealed that things are more intertwined than what we could think (i.e., in several aspects the fields coincide and distinctions are not always so trivial). Therefore discovering analogies and differences may really matter when we want to consider in which way pattern recognition shapes its scientific dimension. Finally we devoted our last contribution to discuss the philosophical underpinnings of much of contemporary pattern recognition research. In the light of Kuhn's notion of a paradigm we discovered that pattern recognition is manifesting signs of a paradigm shift (from a feature-based to a similarity-based approach).

We have briefly summarised our work. But what lesson we learnt from it? What is the take-home message? At first, we learnt that pattern recognition can be genuinely associated to both the practice of science and the practice of technology. In particular we understood that the design of pattern recognition methods could stimulate many philosophical questions (such as: "what is a pattern?" or "what is the relationship between particulars and universals?") as the same way as the research activity does in more traditional scientific disciplines like physics. Then, the philosophical reflection affects basically the implicit or tacit knowledge, what we have called, with Kuhn, “metaphysical paradigms”. Secondly we learnt that if we think of categorization as a relationship between an agent and a set of objects we will be probably inclined to consider pattern recognition and cognitive psychology as the two sides of the same coin. Indeed, our experience suggested that while the field of pattern recognition tends to emphasize “the side of objects” (the feature-based approach in this sense is paradigmatic), the field of cognitive psychology tends to focus more on “the side of the agent objects” (think of, e.g., the employment of empirical observations). In general the emerging points of our research suggest us some future directions. Some of them could contribute to enrich our understanding of pattern recognition and could be in principle developed in the light of some achievements in the philosophy of science. Specifically it could make sense to investigate how the methods built for pattern recognition can be analysed from the standpoints of Lakatos’ “research programmes” and Laudan’s “research traditions,”[55, 59]. Or it could be interesting to extend the examination of Laudan’s philosophy of science, in particular to investigate how the notion of “theoretical problems” [59] affects the research in pattern recognition and machine learning. Another interesting philosophical work could be the disambiguation of some general issues which entered quietly the area of pattern recognition (such as the problem of realism). This activity could be a “natural” continuation of the analysis of the tacit commitments influencing much of pattern recognition research. Moreover we would be really interested also in deepening the interactions between cognitive psychology and pattern recognition. For example it might be interesting to take cue from psychology for the design of experiments in pattern recognition or to study different models of object abstraction even in consideration of the problem of object recognition.

As for the overall research experience we would like to conclude with some considerations on interdisciplinary activity. For sure what emerged from our investigation is that pattern recognition is a variegated research. Sometimes we found entering on the problems of pattern recognition a bit frustrating and discouraging. There were circumstances in which we supposed to understood something that turned out as much interesting as hard to be conveyed. This held especially for philosophy where the abundance of literature and the language of long-standing traditions posed several challenges. And the reason is apparent since the problem underlying the field of pattern recognition may be considered the philosophical problem *par excellence*. The majority of philosophical debates (e.g., realism vs anti-realism, idealism vs empiricism, etc.), indeed, were built around it. On the other hand the the degree of technicality spread across the field of pattern recognition is so impressive that one might be seriously challenged in finding connections to other types of investigation (even to close discipline like cognitive sciences). Moreover the development of sub-communities, built over the success of specific approaches or applications, and accordingly of sub-languages increased the complexity of interdisciplinary projects (for instance, even a work traversing similar areas like computer vision and machine learning is nowadays considered interdisciplinary). Someone could rightly claim that this is the process of over-specialization which, in the end, affected the development of several scientific disciplines. But there are problems which are universal per se and pattern recognition is one of them. Hence, it is not by chance that the term is used in many different disciplines and that, as we suggested in our work, several researches within the field of pattern recognition have posed repeatedly many philosophical concerns. But in such contrasting tendencies (the one towards particular and the one towards universal) what could be most dangerous is the flattening of research, that is thinking that we could easily raise specialized investigation up to more general discussions and, vice versa, apply universal notions to specific domains. This risk recall the Watanabe's genuine proposal of maintaining the terminology "neutral", that is the idea of employing the term "pattern recognition" for both human activities and mechanical simulations [120]. But considering the increasing demand for interdisciplinarity and the acceleration of scientific specialization may terminology be still "neutral"?

Bibliography

[1] The evolution of object categorization and the challenge of image abstraction.

[2] Integrating associative models of supervised and unsupervised categorization.

[3] *Minds Mach.*, 14(4), 2004.

[4] A.Carbone, M. Gromov, and P. Prusinkiewicz. *Pattern Formation in Biology, Vision and Dynamics.* World Scientific, Singapore, 2000.

[5] E. Agazzi. From technique to technology: the role of modern science. *Techné*, 4(2):1–9, 1998.

[6] F. Amigoni, M. Reggiani, and V. Schiaffonati. An insightful comparison between experiments in mobile robotics and in science. *Auton. Robots*, 27(4):313–325, 2009.

[7] J.R. Anderson. *The adaptive character of thought.* Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1990.

[8] J.R. Anderson. The adaptive nature of human categorization. *Psychological Review*, 98:409–429, 1991.

[9] C. Andrieu, N.de Freitas, A. Doucet, and M.I. Jordan. An introduction to mcmc for machine learning. *Machine Learning*, 50(1–2):5–43, 2003.

[10] M. Balcan, A. Blum, and N. Srebro. A theory of learning with similarity functions. *Machine Learning*, 72:89–112, 2008.

[11] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–147, 1987.

[12] C. Bishop. *Pattern Recognition and Machine Learning.* Springer, New York, 2006.

[13] C.M. Bishop. *Neural Networks for Pattern Recognition.* Oxford University Press, New York, 1995.

[14] P. W. Bridgman. *The Logic of Modern Physics.* MacMillan, New York, 1927.

[15] J.M. Buhmann. Information theoretic model validation for clustering. *CoRR*, abs/1006.0375, 2010.

[16] M. Bunge. Technology as applied science. *Technology and culture*, 7(3):329–347, 1966.

[17] H. Cohen and C. Lefebvre. *Handbook of categorization in cognitive science*. Elsevier, 2005.

[18] E.W. Constant. *The Origins of the Turbojet Revolution*. Johns Hopkins University Press, Baltimore, 1980.

[19] A. Cornuéjols and L. Miclet. What is the place of machine learning between pattern recognition and optimization? In *Teaching Machine Learning*, 2008.

[20] J.E. Corter and M.A. Gluck. Explaining basic categories: Feature predictability and information. *Psychological Bulletin*, 111(2):291–303, 1992.

[21] F. Crick. The recent excitement about neural networks. *Nature*, 337:129–132, 1989.

[22] N. Cristianini. On the current paradigm in artificial intelligence. *AICom*, in press.

[23] P.J. Denning. Is computer science science? *Communications of the ACM*, 48(4):27–31, 2005.

[24] H. L. Dreyfus. *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press, Cambridge, Massachusetts, 1992.

[25] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. J. Wiley & Sons, New York, 2000.

[26] R. Duin. Artificial intelligence and pattern recognition. `http://www.37steps.com`. personal communication.

[27] R. Duin. Machine learning and pattern recognition. `http://www.37steps.com`. personal communication.

[28] R. Duin and E. Pekalska. 37 steps blog. `http://www.37steps.com`. personal communication.

[29] R. P. W. Duin, F. Roli, and D. de Ridder. A note on core research issues for statistical pattern recognition. *Pattern Recognition Letters*, 23:493–499, 2002.

[30] R.P.W. Duin. Four scientific approaches to pattern recognition. In A.M. Vossepoel and F.M. Vos, editors, *Fourth Quinquennial Review 1996-2001 Dutch Society for Pattern Recognition and Image Processing*, pages 331–337. Delft, 2001.

[31] R.P.W. Duin and E. Pekalska. The science of pattern recognition: Achievements and perspectives. In W. Duch and J. Mandziuk, editors, *Challenges for Computational Intelligence*, pages 221–259. Springer, Berlin, 2007.

[32] M. Grene (ed.). *Anatomy of Knowledge*. Routledge and Kegan Paul, London, 1969.

[33] P. Feyerabend. *Against Method*. New Left Books, London, 1975.

[34] M. Franssen, G. J. Lokhorst, and I. van de Poel. Philosophy of technology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2010.

[35] A. Freno. Statistical machine learning and the logic of scientific discovery. *IRIS*, 1(2):375–388, 2009.

[36] R. Frigg and S. Hartmann. Models in science. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2012.

[37] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, San Diego, CA, 1990.

[38] R. L. Goldstone and J. Y. Son. Similarity. In K. Holyoak and R. Morrison, editors, *The Cambridge Handbook of Thinking and Reasoning*, pages 13–36. Cambridge University Press, Cambridge, UK, 2005.

[39] G. Gutting. *Paradigms and revolutions: appraisals and applications of Thomas Kuhn's philosophy of science*. University of Notre Dame Press, South Bend, Indiana, 1980.

[40] G. Harman and S. Kulkarni. *Reliable Reasoning: Induction and Statistical Learning Theory (Jean Nicod Lectures)*. The MIT Press, 2007.

[41] J. Hartigan. Introduction. In P. Arabie, L. J. Hubert, and G. de Soete, editors, *Clustering and Classification*, River Edge, NJ, 1996. World Scientific.

[42] S.A. Herbert. *The sciences of the artificial (3rd ed.)*. MIT Press, Cambridge, MA, USA, 1996.

[43] A.G. Hoffmann. *Paradigms of Artificial Intelligence: A Methodological and Computational Analysis*. Springer, Singapore, 1998.

[44] R. A. Hummel and S. W. Zucker. On the foundations of relaxation labeling processes. *IEEE Trans. Pattern Anal. Machine Intell.*, 5:267–287, 1983.

[45] D. W. Jacobs, D. Weinshall, and Y. Gdalyahu. Classification with nonmetric distances: Image retrieval and class representation. *IEEE Trans. Pattern Anal. Machine Intell.*, 22:583–600, 2000.

[46] A.K. Jain and R.C. Dubes. *Algorithms for Clustering Data*. Prentice Hall, 1988.

[47] R. C. Jain and T. O. Binford. Ignorance, myopia, and naiveté in computer vision systems. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53(1):112–117, 1991.

[48] W. James. *The Principles of Psychology*. Harvard University Press, Cambridge, MA, 1983. Originally published in 1890.

[49] F.P. Brooks Jr. The computer scientist as a toolsmith ii. *Commun. ACM*, 39(3):61–68, 1996.

[50] L. N. Kanal. On pattern, categories, and alternate realities. *Pattern Recognition Letters*, 14:241–255, 1993.

[51] J. Kleinberg. An impossibility theorem for clustering. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, pages 446–453. MIT Press, 2002.

[52] W. Krohn, E.T. Layton, and P. Weingart, editors. *The Dynamics of Science and Technology*. Reidel, Dordrecht, Holland, 1978.

[53] T. S. Kuhn. *The Structure of Scientific Revolutions*. The University of Chicago Press, Chicago, 1970. 2nd edition.

[54] T. S. Kuhn. *The Road Since Structure*. The University of Chicago Press, Chicago, 2000.

[55] I. Lakatos. *The Methodology of Scientific Research Programmes*. Cambridge University Press, Cambridge, UK, 1978.

[56] George Lakoff. *Women, Fire and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press, Chicago, 1987.

[57] P. Langley. Machine learning as an experimental science. *Mach Learn*, 3:5–8, 1988.

[58] P. Langley. *Elements of Machine Learning*. Morgan Kaufmann, San Francisco, CA, 1996.

[59] L. Laudan. *Progress and Its Problems: Towards a Theory of Scientific Growth*. University of California Press, Berkeley, CA, 1978.

[60] R. Laudan. Kluwer, Dordrecht, NL, 1984.

[61] H. Margolis. *Patterns, Thinking, and Cognition: A Theory of Judgment*. University of Chicago Press, Chicago, 1987.

[62] J. Maritain. *The Degrees of Knowledge*. Charles Scribner's Sons, New York, 1959.

[63] A.B. Markman and A.R. Otto. Cognitive systems optimize energy rather than information. *Behavioral and Brain Sciences*, 34(4):207–208, 2011.

[64] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt and Co., Inc., New York, NY, USA, 1982.

[65] E. Mayr. *The Growth of Biological Thought*. Harvard University Press, Cambridge, MA, 1982.

[66] W.S. McCulloch and W. Pitts. logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biology*, 5:115–133, 1943.

[67] E. McMullin. Laudan's progress and its problems. *Philosophy of Science*, 46(4):623–644, 1979.

[68] E. Morin, editor. *On complexity*. Hampton Press, Cresskill, NJ, 2008.

[69] M.Pavan and M. Pelillo. Dominant sets and pairwise clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(1):167–172, 2007.

[70] G. L. Murphy and D. L. Medin. The role of theories in conceptual coherence. *Psychological Review*, 92:289–316, 1985.

[71] B. Nicolescu, editor. *Manifesto of transdisciplinarity*. State University of New York Press, Albany, NY, 2002.

[72] N.J. Nilsson. Artificial intelligence: Engineering, science, or slogan? *AI Magazine*, 3(1):2–9, 1981.

[73] R.M. Nosofsky. Choice, similarity, and the context theory of classification. *Journal of experimental psychology. Learning, memory, and cognition*, 10(1):104–114, 1984.

[74] A. Oliva and A. Torralba. The role of context in object recognition. *Trend Cognit. Sci.*, 11:520?–527, 2007.

[75] R. Ortner and H. Leitgeb. Mechanizing induction, 2011.

[76] M. Pavan and M. Pelillo. Dominant sets and hierarchical clustering. In *ICCV*, pages 362–369, 2003.

[77] T. Pavlidis. 36 years on the pattern recognition front: Lecture given at ICPR'2000 in Barcelona, Spain on the occasion of receiving the K.S. Fu prize. *Pattern Recognition Letters*, 24:1–7, 2003.

[78] E. Pekalska and R. Duin. *The Dissimilarity Representation for Pattern Recognition. Foundations and Applications.* World Scientific, Singapore, 2005.

[79] M. Pelillo, editor. *Similarity-Based Pattern Analysis and Recognition.* Springer, London, in press.

[80] M. Pelillo and T. Scantamburlo. How mature is the field of machine learning? In M. Baldoni, C. Baroglio, and G. Boella, editors, *Proceedings of the Thirteenth International Conference on Advances in Artificial Intelligence.* Springer, 2013.

[81] M. Pelillo, T. Scantamburlo, and V. Schiaffonati. Computer science between science and technology: A red herring? In *2nd Int. Conf. on History and Philosophy of Computing*, Paris, France, 2013.

[82] M. Polanyi. *Personal Knowledge. Towards a Post-Cristical Philosophy.* Harper & Row, New York, 1958.

[83] K. Popper. *The Logic of Scientific Discovery.* Basic Books, New York, 1959.

[84] K. R. Popper. *The Open Society and Its Enemies.* Routledge, London, 1945.

[85] K. R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge.* Routledge, London, 1963.

[86] H. Poser. On structural differences between science and engineering. *Techné*, 4(2):81–93, 1998.

[87] E. M. Pothos and T. M. Bailey. Predicting category intuitiveness with the rational model, the simplicity model, and the generalized context model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(4):1062–1080, 2009.

[88] E.M. Pothos, A. Perlman, T.M. Bailey, K. Kurtz, D.J. Edwards, P. Hines, and J.V. McDonnell. Measuring category intuitiveness in unconstrained categorization tasks. *Cognition*, 121:83–100, 2011.

[89] E.M. Pothos and A.J. Wills. *Formal Approaches in Categorization.* Cambridge University Press, Cambridge, UK, 2011.

[90] A. W. Price. *Contextuality in Practical Reason.* Oxford, 2008.

[91] H. Radder. The philosophy of scientific experimentation: a review. *Automated Experimentation*, 1(2):2–9, 2009.

[92] M. D. Resnik. *Mathematics as a Science of Patterns.* Clarendon Press, Oxford, UK, 1997.

[93] R. Rorty. *Philosophy and Social Hope*. Penguin Books, London, UK, 1999.

[94] A. Rosenfeld and H. Wechsler. Pattern recognition: Historical perspective and future directions. *Int. J. Imaging Systems and Technology*, 11(2):101–116, 2000.

[95] P. Rossi. *La Nascita della Scienza Moderna in Europa*. Laterza, Roma, 1997.

[96] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Englewood Cliffs, NJ, 1995.

[97] A. N. Sanborn, T. L. Griffiths, and Daniel J. Navarro. A more rational model of categorization. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, 2006.

[98] A.N. Sanborn, T.L. Griffiths, and D.J. Navarro. Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, 117(4):1144–1167, 2010.

[99] T. Scantamburlo. Il problema della categorizzazione: machine learning vs. psicologia cognitiva. In *IX Convegno dell'Associazione Italiana di Scienze Cognitive*, Istituto di Scienze e Tecnologie della Cognizione del CNR, Roma, 2012.

[100] R.J. Schalkoff. Pattern recognition. In *Wiley Encyclopedia of Computer Science and Engineering*. John Wiley & Sons, Inc., New Jersey, 2008.

[101] V. Schiaffonati and M. Verdicchio. Computing and experiments. *Philosophy & Technology*, pages 1–18, 2013.

[102] J. Schopman. Artificial intelligence and its paradigm. *Zeitschrift fr allgemeine Wissenschaftstheorie*, 17(2):346–352, 1986.

[103] O.G. Selfridge. Pattern recognition and modern computers. In *Proceedings of the March 1-3, 1955, western joint computer conference*, pages 91–94, New York, NY, USA, 1955. ACM.

[104] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T.Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29:93–106, 2008.

[105] J. Serra. Is pattern recognition a physical science? In *Proc. 15th International Conference on Pattern Recognition (ICPR)*, pages 33–40, Barcelona, Spain, 2000.

[106] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, Cambridge, UK, 2004.

[107] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 22(8):888–905, 2000.

[108] H. Skolimowski. The structure of thinking in technology. *Technology and Culture*, 7(3):371–383, 1966.

[109] J. J. Sparkes. Pattern recognition and scientific progress. *Mind*, 81(321):29–41, 1972.

[110] N.P. Suh. *Axiomatic Design: Advances and Applications*. Oxford University Press, New York, Oxford, 2001.

[111] P. Suppes. Models of data. In E. Nagel, P. Suppes, and A. Tarski, editors, *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, pages 252–261, Stanford. Stanford University Press.

[112] P. Suppes. What is a scientific theory? In S. Morgenbesser, editor, *Philosophy of Science Today*, pages 55–67. Basic Books, New York, 1967.

[113] P.R. Thagard. Philosophy and machine learning. *Canadian Journal of Philosophy*, 20(2):261–76, 1990.

[114] A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.

[115] L. Uhr. *Pattern recognition: theory, experiment, computer simulations, and dynamic models of form perception and discovery*. John Wiley & Sons, Inc., New Jersey, 1966.

[116] W.A. Vincenti. *What engineers know and how they know it: analytical studies from aeronautical history*. Johns Hopkins University Press, Baltimore, 1990.

[117] U. von Luxburg and B. Schölkopf. Statistical learning theory: Models, concepts, and results. In M.D. Gabbay, S. Hartmann, and J.H. Woods, editors, *Handbook of the History of Logic*, volume 10, pages 651 – 706. Elsevier North Holland, Amsterdam, Netherlands, 2011.

[118] U. von Luxburg, R. C. Williamson, and I. Guyon. Clustering: Science or art? *JMLR: Workshop and Conference Proceedings*, 27:65–79, 2012.

[119] W.H. Ware. Introduction to session on learning machines. In *Proceedings of the March 1-3, 1955, western joint computer conference*, pages 85–85, New York, NY, USA, 1955. ACM.

[120] S. Watanabe. *Pattern Recognition: Human and Mechanical*. John Wiley & Sons, New York, 1985.

[121] J. Watson. Aristotle's posterior analytics: Ii. induction. *The Philosophical Review*, 13:143–158, 1904.

[122] J. Williamson. A dynamic interaction between machine learning and the philosophy of science. *Minds and Machines*, 14(4):539–549, 2004.

[123] A.J. Wills and E. M. Pothos. On the adequacy of current empirical evaluations of formal models of categorization. *Psychological bulletin*, 138(1):102–125, 2012.

[124] D. Wojick. The structure of technological revolutions. In G. Bugliarello and D.B. Doner, editors, *The History and Philosophy of Technology*, pages 238–247. University of Illinois Press, Chicago, 1979.