

# Vision-based Robot Navigation



# Vision-based Robot Navigation

## Quest for Intelligent Approaches

### Using a Sparse Distributed Memory

Mateus Mendes, A. Paulo Coimbra, Manuel M. Crisóstomo



Universal Publishers  
Boca Raton

*Vision-based Robot Navigation*  
*Quest for Intelligent Approaches Using a Sparse Distributed Memory*

Copyright ©2012 Mateus Mendes, A. Paulo Coimbra, Manuel M. Crisóstomo  
All rights reserved.

No part of this book may be reproduced or transmitted in any form or by any means,  
electronic or mechanical, including photocopying, recording, or by any information storage  
and retrieval system, without written permission from the publisher.

Universal-Publishers  
Boca Raton, Florida  
USA • 2012

ISBN-10: 1-61233-104-1  
ISBN-13: 978-1-61233-104-1

[www.universal-publishers.com](http://www.universal-publishers.com)

Cover image @Cutcaster.com/Miro Kovacevic

Library of Congress Cataloging-in-Publication Data

Mendes, Mateus, 1978-

Vision-based robot navigation : quest for intelligent approaches using a sparse distributed  
memory / Mateus Mendes, A. Paulo Coimbra, Manuel M. Crisóstomo.

p. cm.

Includes bibliographical references.

ISBN-13: 978-1-61233-104-1 (pbk. : alk. paper)

ISBN-10: 1-61233-104-1 (pbk. : alk. paper)

1. Robot vision.
2. Mobile robots.
3. Intelligent control systems. I. Coimbra, A. Paulo, 1962- II. Crisóstomo, Manuel M., 1952- III. Title.

TJ211.3.M46 2012

629.8'92-dc23

2012014892

*To Ísis  
Mateus Mendes*

*To Mafalda, Gonçalo and my parents  
A. Paulo Coimbra*

*To Joana  
Manuel M. Crisóstomo*



*Thanks to  
Higher School of Technology and Management of Oliveira do Hospital, IPC, Portugal,  
Institute of Systems and Robotics – University of Coimbra, Portugal,  
and to Fundação para a Ciência e Tecnologia, Portugal,  
for sponsoring part of this work.*





## Preface

Intelligent machines have been in the imaginary of humans, not for decades, but for millennia. Even ancient civilisations dreamt of intelligent machines who could talk, run or protect property and humans from strangers. Nonetheless, for many centuries, those machines were nothing more than fiction. Despite important scientific advances that are still in the basis of modern computers and robots, significant results and attempts at building intelligent machines have only been reported in the 20<sup>th</sup> century. It was only a few decades ago that sensing and motion technologies and programmable computers boosted the area.

The main problem that has to be faced is that of building robots that know where they are, where they have been, and how to go some place they are supposed to know or to be able to find a way to. That is the problem of autonomous robot localisation and navigation. This book reviews the origins of intelligent machines, modern theories on how the brain works and robot localisation and navigation methods. It is focused on a method of navigating robots using visual memories stored into an associative memory. The reason for the choice is based on two important observations: i) vision is, for the average human, the most important source of information; and ii) modern research is showing that the use of a sophisticated memory is perhaps the source of human intelligence.

The associative memory used is a Sparse Distributed Memory (SDM), a kind of associative memory proposed by Pentti Kanerva, based on the properties of high dimensional binary spaces. Some characteristics of SDMs are their high immunity to noisy data, ability to deal with incomplete information and learn on a single pass, and also to *naturally* forget older memories when close to their maximum capacity. They work very well on high dimensional spaces, and are also appropriate to work with sequences of events, where each event is associated with its successor or predecessor. The SDM naturally confers on the robot many characteristics of the human memory, thus being a powerful and general model, comparable to or better than other navigation methods.

Since the SDM was proposed in the 1980s, it has been subject to intense research and many authors have proposed improvements. However, so far it was not clear how robust the SDM would be in practise, specially in the task of robot navigation, compared to other modern approaches.

The present book intends to be a useful tool for students and researchers working towards building intelligent machines. The first chapter intends to give the reader an idea of the historical background behind intelligent machines. Chapter 2 compares human and artificial intelligence. It briefly reviews the most up to date definitions of intelligence, and briefly describes two expected sources of intelligent behaviour: the human brain, which is thought to be the source of human intelligence, and the Turing machine, which is able to run any computable algorithm, and so is the natural candidate to use to try to build intelligent

---

machines. This chapter also briefly reviews the state of the art in terms of associative memory models.

Chapter 3 is devoted to the Sparse Distributed Memory model, as proposed by Pentti Kanerva and later improved by many different authors. It summarises the mathematical principles behind the SDM and Kanerva’s original proposal to build the memory.

Chapter 4 is a review of the state of the art in terms of robot navigation approaches. It describes the more important paradigms that have been used in recent years to map and navigate robots, focusing specially approaches that are based on the use of visual memories.

Chapter 5 describes an example application where a robot is guided by visual memories stored into a SDM. It also describes the different operation modes used, as well as the performance assessment measures that were defined to compare the results obtained using the different operation modes. It ends with the conclusion that the performance of the original SDM model is impaired because it is based on the use of a Hamming distance, which does not consider the position of the bits, while the sensory information is based on the use of the natural binary code, in which the value of each 1 depends on its position. This problem is called the “encoding problem.”

Chapter 6 is dedicated to the analysis of the encoding problem. The problem arises from using different criteria to encode the information that comes from the sensors and to process it inside the SDM. Therefore, it can be solved either by encoding the sensory information in a different way, or by changing the structure of the SDM to work differently.

Chapter 7 is a study of the performance of the SDM under typical robot navigation problems. Problems such as that of the kidnapped robot, illumination changes and scenario changes, memory overflow and forgetting are debated in detail.

Chapter 8 describes a method of navigating a robot using visual memories. The SDM is used both to store images and as a pattern matching tool to build a topological map.

Chapter 9 summarises all the contents of the book and the main advantages and disadvantages of the approach described. It also opens perspectives of possible future work to further research in these topics.

Finally, Appendix A describes a comparison of the performance of the SDM in two different domains: the domain of robot navigation using view sequences, and that of robot manipulation using position sensors. The former domain is characterised by long vectors and large amounts of noise that are found in the images. The latter domain is characterised by relatively short vectors and very small amounts of noise. The appendix describes experiments and compares the results obtained.

# Contents

<b>1. The Quest for Artificial Intelligence</b>	<b>1</b>
1.1. A Dream Thousands of Years Old	1
1.1.1. The Origins	1
1.1.2. Modern History	3
1.2. The Problems	6
1.2.1. Poorly Defined Problem, or Wrongly Defined Problem	6
1.2.2. High Expectations	8
1.2.3. Environmental Difficulties	10
1.3. Summary	10
<b>2. Intelligence and the Brain</b>	<b>13</b>
2.1. Human and Artificial Intelligence	13
2.1.1. The Turing Machine	14
2.1.2. Definitions of Intelligence	15
2.2. On the Brain	18
2.2.1. Brain Structure	19
2.2.2. The Brain as a Memory System	21
2.2.3. The Memory-Prediction Framework	22
2.2.4. Invariant Forms and Hierarchical Temporal Memory	24
2.2.5. A Critical View of Hawkins' Work	26
2.3. Some Memory Models	29
2.3.1. Associative Memories	29
2.3.2. Willshaw Network	29
2.3.3. Hopfield Network	31
2.4. Summary	32
<b>3. Sparse Distributed Memory</b>	<b>33</b>
3.1. Introduction	34
3.2. The space $2^n$	35
3.2.1. Concepts and relations	35
3.2.2. Properties of $2^n$	37
3.2.3. Memory Items As Points of $2^n$	40
3.3. A Sparse, Distributed Memory	41
3.3.1. Memory Size and Sparseness	41
3.3.2. Writing and Reading	42
3.4. Implementation of a SDM	43
3.4.1. SDM Using Linked Lists	44

3.4.2.	SDM Using a Neural Network . . . . .	46
3.5.	Characteristics of the SDM . . . . .	47
3.6.	Previous Applications and Improvements . . . . .	49
3.6.1.	The Stanford Model . . . . .	49
3.6.2.	Weather Forecasting . . . . .	50
3.6.3.	Object Indexing . . . . .	51
3.6.4.	Jaeckel's Selected-Coordinate Design . . . . .	52
3.6.5.	Hely's Signal Model . . . . .	54
3.6.6.	On-Line Value-Based Reinforcement Learning . . . . .	55
3.6.7.	SDM Using M-of-N Codes . . . . .	56
3.6.8.	Working with Sequences . . . . .	58
3.6.9.	Image Recognition . . . . .	61
3.7.	SDM Advantages and Disadvantages . . . . .	61
3.8.	Summary . . . . .	63
<b>4.</b>	<b>Robots Localisation and Navigation</b>	<b>65</b>
4.1.	Pipeline and Subsumption Architectures . . . . .	65
4.1.1.	Pipeline Architecture . . . . .	66
4.1.2.	Subsumption Architecture . . . . .	66
4.2.	Mapping and Navigation Paradigms . . . . .	68
4.2.1.	Occupancy Grids . . . . .	69
4.2.2.	Potential Fields . . . . .	70
4.2.3.	Voronoi Diagrams . . . . .	71
4.2.4.	Topological Maps . . . . .	72
4.3.	View-Based Navigation . . . . .	75
4.3.1.	The View-Based Paradigm . . . . .	75
4.3.2.	Matsumoto's Approach . . . . .	77
4.3.3.	Ishiguro's Approach . . . . .	79
4.3.4.	The T-Net . . . . .	79
4.3.5.	Mallot's Approach . . . . .	80
4.3.6.	Other View-Based Approaches . . . . .	82
4.3.7.	Associative Memory Using a Look-up Table . . . . .	82
4.4.	Previous Work Using SDMs . . . . .	83
4.4.1.	Rao and Fuentes Goal-Directed Navigation . . . . .	83
4.4.2.	Watanabe's AGV . . . . .	85
4.4.3.	Similar Techniques . . . . .	85
4.5.	Summary . . . . .	89
<b>5.</b>	<b>Vision and Memory-Based Robot Navigation</b>	<b>91</b>
5.1.	Experimental Platform . . . . .	91
5.1.1.	Hardware . . . . .	92
5.1.2.	Testbed . . . . .	93
5.2.	Control and Image Enhancement Software . . . . .	94

5.2.1.	Operational Level . . . . .	94
5.2.2.	Image Processing . . . . .	96
5.2.3.	Focus Detection by Contrast Measurement . . . . .	101
5.3.	SDM Implementation . . . . .	104
5.3.1.	First Design . . . . .	104
5.3.2.	Bitwise Implementation . . . . .	106
5.3.3.	Arithmetic Implementation . . . . .	108
5.3.4.	Noisy Data and Use of Dynamic Access Radius . . . . .	109
5.4.	Implementing Vision-Based Navigation . . . . .	113
5.4.1.	Navigation Process . . . . .	113
5.4.2.	Experiments and Preliminary Results . . . . .	116
5.5.	Arithmetic and Hamming Distances . . . . .	120
5.5.1.	Comparing Distances . . . . .	120
5.5.2.	Comparing Mapping Functions . . . . .	121
5.5.3.	Momentary Localisation Errors . . . . .	123
5.6.	Summary . . . . .	124
<b>6.</b>	<b>Data Encoding Issues</b>	<b>127</b>
6.1.	Binary Codes and the Hamming Distance . . . . .	127
6.1.1.	The Problem of Using the Hamming Distance . . . . .	128
6.1.2.	Gray Code . . . . .	130
6.1.3.	Sorting the Bytes . . . . .	131
6.1.4.	Reducing Grey Levels . . . . .	133
6.2.	Experiments and Results . . . . .	135
6.2.1.	Experiments without Data Distribution . . . . .	135
6.2.2.	Experiments with Data Distribution . . . . .	140
6.3.	Summary . . . . .	144
<b>7.</b>	<b>Memory Performance</b>	<b>145</b>
7.1.	The Kidnapped Robot Problem . . . . .	145
7.2.	Robustness to Noise . . . . .	148
7.2.1.	Illumination Changes . . . . .	148
7.2.2.	Scenario Changes . . . . .	150
7.3.	Reinforcement of an Idea . . . . .	155
7.4.	Memory Saturation . . . . .	156
7.4.1.	The Problem of Saturation . . . . .	156
7.4.2.	Sequence Disambiguation . . . . .	159
7.4.3.	Memory Overflow . . . . .	166
7.5.	Retrieving Sequences in Reverse Order . . . . .	169
7.6.	Outdoor Experiments . . . . .	171
7.7.	Summary . . . . .	173
<b>8.</b>	<b>Navigation Using View Sequences</b>	<b>177</b>

8.1.	The Problem of Path Planning . . . . .	177
8.1.1.	The Problem . . . . .	178
8.1.2.	Building a Topological Map . . . . .	178
8.1.3.	Goal Directed Navigation . . . . .	180
8.2.	Navigation Experiments and Results . . . . .	181
8.2.1.	Experiments in Reconstitution of Country Environment . . . . .	182
8.2.2.	Experiments in Reconstitution of Urban Environment . . . . .	182
8.3.	Improved Performance Using a Sliding Window . . . . .	183
8.3.1.	Distribution of the Momentary Localisation Errors . . . . .	185
8.3.2.	Use of a Sliding Window . . . . .	186
8.3.3.	Experiments and Results . . . . .	186
8.4.	Possible Applications . . . . .	188
8.5.	Summary . . . . .	191
<b>9.</b>	<b>Conclusions</b>	<b>193</b>
9.1.	Overview of the Topics Covered . . . . .	193
9.1.1.	What Intelligence Is . . . . .	193
9.1.2.	Brain Models . . . . .	194
9.1.3.	Intelligent Robot Navigation Using View Sequences . . . . .	195
9.1.4.	The Encoding Problem . . . . .	195
9.1.5.	Limits of the SDM . . . . .	196
9.1.6.	Path Planning . . . . .	196
9.2.	Advantages and Disadvantages of the Approach Described . . . . .	197
9.3.	Open Research Questions . . . . .	198
<b>A.</b>	<b>Cross-Modal Comparison</b>	<b>201</b>
A.1.	Introduction . . . . .	201
A.2.	TAMS Service Robot . . . . .	202
A.2.1.	Hardware . . . . .	202
A.2.2.	Arm Control . . . . .	204
A.3.	Experiments and Results . . . . .	204
A.4.	Discussion . . . . .	205
A.5.	Summary . . . . .	207
	<b>Acronyms and Symbols</b>	<b>209</b>
	<b>List of Figures</b>	<b>210</b>
	<b>List of Tables</b>	<b>213</b>
	<b>References</b>	<b>215</b>

# 1. The Quest for Artificial Intelligence

*If you have an apple and I have an apple and we exchange apples then you and I will still each have one apple. But if you have an idea and I have one idea and we exchange these ideas, then each of us will have two ideas.*  
(George Bernard Shaw, Irish dramatist)

## Contents

---

<b>1.1. A Dream Thousands of Years Old</b> . . . . .	<b>1</b>
1.1.1. The Origins . . . . .	1
1.1.2. Modern History . . . . .	3
<b>1.2. The Problems</b> . . . . .	<b>6</b>
1.2.1. Poorly Defined Problem, or Wrongly Defined Problem . . . . .	6
1.2.2. High Expectations . . . . .	8
1.2.3. Environmental Difficulties . . . . .	10
<b>1.3. Summary</b> . . . . .	<b>10</b>

---

The ultimate goal of robotics is to build intelligent machines, fit for a myriad of different jobs. In the last fifty years, the emergence of a new research area, called Artificial Intelligence (AI), shed new light and opened new perspectives to the field [13, 91, 80]. This chapter briefly reviews some important historic aspects and the main problems encountered during the quest towards intelligent machines. It also gives an outline of the book structure.

## 1.1. A Dream Thousands of Years Old

Robotics and Artificial Intelligence are usually regarded as modern research areas. Indeed, the term “Artificial Intelligence” was coined just in the 20<sup>th</sup> century, as the main subject of a conference that was held in Dartmouth in 1956 [68]. But the key ideas of artificial beings can be traced back thousands of years [13, 91, 80]. This section briefly reviews some important historic landmarks, from ancient to modern history, more related to the theme of intelligent robot navigation.

### 1.1.1. The Origins

The modern word “computer” was derived from an old human job. The concept was adjusted to be applicable to machines. But it is not only the word that is hundreds of years old. The

dream of building intelligent machines, and some important contributions to the field, date back to at least the ancient Greeks, Egyptians and Chinese.

### The word computer

It is a popular saying that “while humans do not build intelligent robots, humans are the robots.” And that was true with computers in a literal sense. The term “computer” was used to refer to a human job that consisted in performing routine calculations on a daily basis. Those human computers spent days calculating navigational tables, tide charts, planetary positions for astronomical almanacs and other tasks which are currently performed by modern digital computers. However, humans get bored and tired with ease, and inevitably start making more and more mistakes as time goes by. Besides, they also take their time—usually more than a machine for that kind of jobs. As soon as the technology was developed, machines were built to make those routine tasks automatically. The process was sped up specially during the Second World War, when top scientists were called to break German military’s cipher [41]. Those first machines were called “automatic computers,” so that they could be differentiated from the “human computers.” That started happening as early as the 1920s [81].

### Credits to ancient civilisations

The quest for intelligent machines is even older than the creation of the modern computers. As early as 2500 B.C., the Egyptians developed the concept of “thinking machines,” which were regarded with some mysticism. There are reports that back in the ancient Egypt, 800 B.C., a statue was built that could move its arm and speak to onlookers. However, its “intelligence” was that of a man that could fit inside the husk [34]. Nonetheless, those statues that could move their limbs were the first automata, which are the predecessors of the modern robots.

The concept of robots is also very present in the Greek mythology. According to the legends, Hephaestus, the God of the smiths, discovered ways of working iron, copper and other metals. Using that expertise, he built mechanical servants to serve him<sup>1</sup>. One of those amazing creatures made by Hephaestus was Talos<sub>1</sub>. Talos<sub>1</sub> was a robot made of bronze, who guarded Crete by running around the island three times every day. When intruders appeared at the island, he pelted them with stones. There are many different versions of the legend, but the most interesting is perhaps that in which Talos<sub>1</sub> is a *man*, the last of a generation of *men* made of bronze [93].

Besides the mythological concepts, another important contribution to building intelligent machines took place in the ancient Greece. The Greek philosopher Aristotle, in the 5<sup>th</sup> century B.C., created the “syllogistic logic.” The syllogistic logic is the origin of deductive reasoning and the basis for implementing intelligent behaviours in artificial agents.

---

<sup>1</sup><http://homepage.mac.com/cparada/GML/Hephaestus.html> (last checked 2010.03.16).



Those concepts of robots and intelligent reasoning had to be combined with some technology in order to produce the first machines (and even the most advanced modern robots to date are not even close to the mythic Talos<sub>1</sub>). The oldest ingenious machine is possibly the abacus. It is not clear when the first abacus was created, but it is usually credited to the Babylonians, who might have developed the first abacus circa 2400 B.C. [97]. The abacus is just an auxiliary memory to the human users who have to perform the calculations. It helps and speeds up the process of making basic arithmetic operations, namely sums and subtractions. Division and multiplication are more complicated to perform using abaci.

There are also reports of an ingenious device, called the “South Pointing Chariot,” which may have been invented in China as early as 2600 B.C [19]. The South Pointing Chariot was possibly the first mechanical device that used differential gearing. The mechanism was built in a way that a figure (figure of a hand, dragon, etc.) atop the chariot maintains its heading, always pointing the South, regardless of the path travelled by the chariot. Different versions of the device have been built. The worst versions accumulated a large error in just a few turns. The best versions were able to handle hundreds of turns with just a minimum error, thus being useful navigation devices for those who could afford such an equipment.

### 1.1.2. Modern History

Modern history can be described as an exponential line of successes in various sub-fields. Nonetheless, it is made of ebbs and flows in each sub-field. From the first mechanical calculator built in the 17<sup>th</sup> century to the modern humanoid robots that can play soccer [7], there have been periods of great optimism, but also “winters” of doubt which were only broken when there was a new breakthrough.

#### The first calculator

Many centuries passed since the first attempts to build intelligent machines. Developments in mathematics and physics were very slow, but in the 17<sup>th</sup> century the first mechanical calculators were built. The Pascalina<sup>2</sup> (see Figure 1.1), built by Blaise Pascal, was the first commercial calculator. It had only a very limited ability to sum and subtract. Therefore, it was just a more sophisticated abacus, based on gears instead of movable stones. During his youth, Pascal built about 50 prototypes and sold about 12 units of those Pascalinas, before losing interest in the commercialisation of the device.

#### Robots and Artificial Intelligence

The term “robot” was used for the first time by Karel Capek, in a science fiction novel published in 1921 in the Czech Republic [15]. Capek envisioned machines created as mechanical slaves to work for human masters, and called those machines “robots.” In the novel, the

---

<sup>2</sup><http://www.computernostalgia.net/articles/pascalina.htm> (last checked 2010.03.16).

<sup>3</sup><http://www.sciencemuseum.org.uk/images/I032/I0302630.aspx> (last checked 2010.03.16).



Figure 1.1.: Picture of a replica of the first commercial calculator, which is in exhibition at the Science Museum in London<sup>3</sup>.

robots are treated poorly, but they are “happy” until they are programmed to incorporate emotions. At that point, the robots rebel in order to take over the world and get free from their human masters.

Out of the science fiction world, many innovations took place in the early years of the 20<sup>th</sup> century. The first and second World Wars sped up the technological developments. The first modern computer was built in the 1940s, mostly with technology developed during the war time. The end of the Second World War was actually a period of great optimism, in what concerns technological development. The technologies developed in the military industry were released for civil applications, for the benefit of the public in general. Science was making widely available artifacts which a few decades before were just visions of science fiction writers. Therefore, almost everything seemed possible to accomplish in a very short time frame. It was expected that in a few decades intelligent robots would be everywhere, performing not only repetitive computing tasks, but also tasks which required high-level reasoning, learning and ability to perform intelligent decisions, as well as other skills which usually can only be achieved by human workers. In the 1950s many authors were optimistic enough to predict that in 10–20 years computers would be the world chess champions<sup>4</sup> and uncover new mathematical principia [91].

The term “Artificial Intelligence” (AI) was coined in 1955, by John McCarthy [68]. AI was meant to be the subject of a scientific summit in Dartmouth. That summit brought together some of the most important researchers at that time, in order to discuss the new scientific and technological developments. The subject of the meeting actually became the name of a whole new area of research.

---

<sup>4</sup>Actually, it was only in 1997 that the supercomputer Deep Blue was able to beat the world chess champion Garry Kasparov. However, Deep Blue was not a mere “general purpose” computer. Deep Blue was a supercomputer designed specifically for the purpose of defeating Kasparov.

## **Ebbs and flows**

The initial optimism and faith in the achievements of science and technology encountered unexpected barriers, when the ideas were to be put into practise. Apparently simple tasks turned out to be incredibly complicated to perform using robots or computers. The first huge problem spotted was that computers were actually performing simple algorithms, with little knowledge of their own about the context. They had no *awareness* and no *common sense*, which seem to be basic human skills that guide many decisions and thought processes. That lack of *awareness* is in the origin of the still open question launched by John Searle later in the 1980s, which is discussed in more detail in Chapter 2.1.2.

In the 1970s AI has gone through what is known as the first “AI winter”—about six years of disappointment with the results obtained and extensive cuts in research funds. In the early 1980s, the rise of expert systems caused another wave of optimism. Expert systems are simple computer programs that, using logical rules which are inferred from the knowledge of human experts, mimic their actions and answer questions correctly. That approach was very welcome by companies, which could then use machines to replace human experts, saving money and reducing the risk of human failure and boredom. In the late 1980s, though, the enthusiasm faded again, in part due to the realisation that expert systems are not a general solution to all kind of reasoning problems. Expert systems are only fit for a small number of decision-making applications. By that time, a new debate arose, with many authors questioning if intelligence was possible without a body. That question is in the origins of the “Behaviour-Based Artificial Intelligence” (BBAI), which builds on the idea that intelligence is not possible without a body. BBAI is tightly related to Behaviour-Based Robotics (BBR), which builds on the same idea. According to BBR, intelligent behaviour arises from a number of small modules, which react to stimuli in a semi-autonomous way. A popular approach is the subsumption architecture, proposed by Rodney Brooks, which is explained in more detail in Chapter 4.1.

## **How far can machines evolve?**

Robotics and AI have known many successes. Numerous techniques, such as expert systems, neural networks, data mining and many others, have found large-scale commercial applications. However, up to this day, few of the initial promises of AI have been fulfilled. The debate is still open, on how intelligent will be the most intelligent artificial machines, or the ultimate robot. The optimists say that it is just a matter of time. With increasing computer power and memory, the problem is just to put everything together. The pessimists say that artificial machines will have to be radically different from the ones currently built if they are meant to reach levels of intelligence and competence comparable to that of human beings, because it is not possible to capture the true essence of intelligence using an algorithmic approach. The problem may be that research is following the wrong track: intelligence is the result of a process radically different from computing. Claude Shannon’s discussion on “What Computers Should Be Doing” is often quoted in this aspect [16]. Shannon states:

Efficient machines for such problems as pattern recognition, language translation, and so on, may require a different type of computer than any we have today. It is my feeling that this computer will be so organised that single components do not carry out simple, easily described functions. ... Can we design... a computer whose natural operation is in terms of patterns, concepts, and vague similarities, rather than sequential operations on ten-digit numbers?

If the pessimists are right, current research may be to a great extent on the wrong path. However, the history of Artificial Intelligence and Intelligent Robots has been made of ebbs and flows, much like the economic cycles. Occasionally, a new discovery spikes a wave of optimism, which soon fades and gives room to a “winter” of pessimism and only small advances. The current sentiment among researchers is that new approaches are still needed to renew AI, if the goal of building intelligent robots is to be achieved in the mid term. One of the most ambitious projects is the Mind Machine Project, launched in 2009. It is a 7-year and multi-million dollar project, involving some of the most renowned researchers in the field<sup>5</sup>. The Mind Machine Project intends to rethink AI from scratch, including fundamental assumptions such as the nature of the mind and memory. That is one major step towards maintaining AI’s original goal of building machines of human-level intelligence, while many authors agree that the current research is mostly split into a wide range of sub-areas. In those sub-areas, the focus is on building intelligent machines for specific tasks where a limited level of intelligence is needed, without caring about the big picture or cognitive insight into the nature of the problems<sup>6</sup>. The Mind Machine Project intends to shed new light onto the problem, using a holistic approach.

## 1.2. The Problems

As explained in Section 1.1, the dream of creating intelligent machines is thousands of years old. However, the most significant developments in the field date back only a few tens of years, and it is still not clear if the goal can be achieved at all. The difficulties start with the very definition of the problem, which is itself very unclear. Besides, there are very high expectations, and those expectations are raised continuously. However, the knowledge and technologies available are still rudimentary, and may even be insufficient to achieve the final goal.

### 1.2.1. Poorly Defined Problem, or Wrongly Defined Problem

The very definition of the goal is not clear at all. What is intelligence? Is it unique, or are there different levels and types of intelligence? What types of intelligent behaviours can be mimicked and are suitable to implement in artificial beings? What behaviours can be programmed in modern computers?

---

<sup>5</sup>For more information on the MIT Mind Machine Project see <http://mmp.mit.edu> (last checked 2010.03.10).

<sup>6</sup>See, for example, [http://www.itweb.co.za/index.php?option=com\\_content&view=article&id=30637](http://www.itweb.co.za/index.php?option=com_content&view=article&id=30637) for a more complete discussion (last checked 2010.03.18).

## The goal

The first problem that is encountered when the goal is to build intelligent machines is that of understanding what intelligence is. Chapter 2 discusses the problem in detail. The definition of intelligence may just encompass the success of the intelligent agent, or it may also care about *how* the agent does it. In the first case it is only the performance that matters. The latter case is more complicated, for it involves other concepts which are also poorly defined. The agent is expected to *understand* or *comprehend* its actions and its environment. However, there are no clear definitions of understanding and comprehending, as explained in Chapter 2.

Hence, the problem in its general formulation is poorly defined and comprises a goal which is not clear. Additionally, there is no standard way to measure the intelligence of a machine or a human being. Of course, there are the Intelligence Quotient (IQ) tests, but those are still subject to many criticisms and are not a standard measure of intelligence ready to be applied to artificial minds—and even more so when the goal is just to implement *intelligent* robot navigation.

## Is the goal reachable?

Actually, the problem of reasoning and making *intelligent* decisions in general may not be solvable at all using currently known technology, as suggested in Section 1.1.2. For example, it is known that a sentence or word may cause the release of a chemical substance in the human brain. That chemical substance may then cause a feeling or a change into a mood. Then, after the sentence is forgotten, the feeling continues [50] until the chemical substance that causes it fades. This means that in the human brain the feeling is separated from thought. However, the feelings influence behaviours and decisions. That aspect of the human way of action is completely alien to machines, which are built upon a completely different *anatomy*. Even the most modern machines are highly (blindly) rational and algorithmic, while humans rely on a complex amalgam of emotions, memories and sensations to judge and act. Of course, even the release of a chemical substance and its fading can be simulated in software, although it is an awkward way to program decision-making processes. Nonetheless, signal-propagation models, for example, may be adequate to simulate those behaviours.

Another important aspect is that the brain seems to have a natural ability to do mind reading. It is now well known that humans in general are very keen on mind reading, and that ability plays an important role in normal human communication and social relations. In a normal conversation, most of the information is exchanged non-verbally, and even unconsciously. Humans easily spot if someone is sad or happy, confident or unsure, nervous or relaxed. A simple look at a face, even at a picture of a face, is enough for the brain to unconsciously interpret a myriad of signals and spot a lot of emotions. Those non-verbal and unconscious analysis also play an important role when two people do not share a common language, or when someone is cheating or lying. The ability is natural in humans and to a great extent independent from race and culture. Paul Ekman and Tomkins [31] identified

a set of about 300 facial expressions that are responsible for the ability to express about 30 different emotions. Those basic expressions are called Action Units (AU). An AU is, e.g., narrow lips or nose. Those 300 AU combined make around 10000 useful combinations, allowing humans to express emotions such as anger, happiness and disgust, among others. That ability plays a very important role in human communication—sometimes even more important than verbal communication. Computers are in a certain way like autists. There is a long way to go to program them to express non-verbal language and also to understand non-verbal signals. Computers will surely still lack the abilities to read someone’s mind and to express non-verbal signals for quite a long time. The ability to automatically read emotion from a facial expression may not be impossible to program, but it seems to be very far away. The most promising approach might be using sophisticated computer vision algorithms, able to distinguish an AU in an image of a face. Actuators strategically placed in an artificial skin might provide the necessary flexibility to imitate such facial expressions. Such abilities will make it possible to build a robot with very good conversation skills<sup>7</sup>.

### **Paving a new way**

An important consideration is that imitating humans may not be important at all to build intelligent machines. It is often said that men did not learn to fly by imitating birds. The human solution to the problem of flying is completely different of that used by the birds and insects that fly. Yet, the human way is suitable for human needs, safe and efficient<sup>8</sup>. Hence, it may happen with robots the same that succeeded with the human ability to fly. A completely different approach may be developed, resulting in the construction of machines which are different from humans, but equally smart and fit for many different purposes. The result may be a new, non biological, form of life. Current researchers tend to either imitate nature, being strongly tied to the natural models, or to neglect it completely, not taking advantage of what can be learnt from the natural world. But many different approaches may be valid to achieve the goal of building intelligent machines, and so far it is unclear which one is the more promising or more suitable one.

#### **1.2.2. High Expectations**

There have always been high expectations, when the theme is “intelligent machines.” Machines have to be smart, at least as much as humans, but cannot fail. And when a particular

---

<sup>7</sup>See for example the Felix project at <http://www.felix-growing.org/about> (last checked 2010.03.30). Felix pursues the goal described.

<sup>8</sup>Despite the success of the flying devices that use different approaches, the construction of machines able to fly by flapping wings, like birds, is more energy efficient, and it is still an open area of research. The website [www.ornithopter.org](http://www.ornithopter.org) contains a thorough compilation of information about such devices, called “ornithopters.” Todd Reichert was the first human to fly on an ornithopter, on August 2, 2010. That ornithopter had wings that spanned for 32 meters, and it was made of light materials, such as carbon fibre, foam and balsa wood. It weighed only 43 Kg and was moved by pedalling. Todd Reichert flew the aircraft for 19.3 seconds. (Source: <http://www.ottawacitizen.com/dreams+flies/3581043/story.html>, last checked 2010.09.26.)

sub-goal is reached, it loses its mysticism and is not considered a sign of intelligence any longer.

## **Better than humans**

From the ancient Egyptians to the Dartmouth summit, and now to the modern think tank that is the Mind Machine Project, the dream has always been the same: to build intelligent machines, that can think and act like humans. However, the fact is that those intelligent machines have to be better than *average* human beings. Everyone is, to a great extent, complacent with an occasional, even if unexpected and apparently unexplained, human failure. The Roman philosopher Seneca wrote the famous sentence “errare humanum est” (to err is human), and it is still valid and popular nowadays. But an intelligent machine is not expected to make a mistake, under “normal” circumstances. And, depending on the kind of machine, a machine error may endanger people or property, perhaps more than a human failure in general. Therefore, researchers can hardly be satisfied with a machine of human-like abilities. Actually, the goal is to produce biologically inspired machines, but flawless in the range of problems they are designed to solve and the jobs they are made to accomplish.

## **A moving target**

Another problem is that Artificial Intelligence is actually pursuing a kind of “moving target.” The concept of intelligence is poorly defined. Hence, it is hard to classify any system as intelligent or unintelligent. A popular example is ELIZA, the first chatter bot, written in 1966 [108]. ELIZA is able to maintain a conversation with a human in a limited context, the most famous of which is the emulation of a Rogerian psychotherapist<sup>9</sup>. Many people mistook ELIZA for a real person and attributed it superb intelligence. However, ELIZA is just the implementation of a few clever tricks. The key idea is to pick some selected keywords in the user’s sentence and reply with a selected question based on that keywords. For example, to the sentence “I’m sad,” a very probable (and acceptable) ELIZA answer is “How long have you been sad?” The keyword “sad” triggers the question. Other keywords trigger other questions which make some sense because they are on-topic. That creates an illusion of intelligence, while in fact ELIZA does not even have a memory or context of the conversation. It works solely based on pattern matching. The question, which has been largely debated, still remains: is it intelligent? Despite the initial enthusiasm, nowadays it is widely accepted that intelligence has to be much more than those simple ELIZA-like tricks. The subject is discussed in more detail in Section 2.1.2.

---

<sup>9</sup>To test a working implementation of ELIZA see, for example, the application online at <http://www.manifestation.com/neurotoys/eliza.php3> (last checked 2010.03.21).

### 1.2.3. Environmental Difficulties

As presented in Section 1.2.1 (and discussed in more detail in Section 2.1.2), the concept of intelligence is poorly defined. Therefore, researchers are trying to find their way along unknown terrain, without a clear understanding of the localisation of the goal or the best way to find it. And there is also another problem: the environments that machines have to face can be quite hostile and complex, while the robots in general still possess only poor sensing abilities. Russell & Norvig classify the environments according to the following five properties [91]:

- **Accessible/inaccessible**—Accessible is that environment which can be fully sensed by the robot. The robot must have access to all the information that is relevant to its choices and possible behaviours. The robots usually have limited sensory information, so they have to work with an only partially accessible environment.
- **Deterministic/nondeterministic**—In a deterministic environment, the next state is completely determined by the present state and the robot's actions. Real environments are, therefore, nondeterministic, from the point of view of the intelligent agent. Environments with inaccessible properties are always nondeterministic, for the agent cannot determine exactly the state of the variables it cannot sense.
- **Episodic/nonepisodic**—In episodic environments, the agent's experiences can be divided into sequences of perceptions and acts, which are self-contained episodes. What happens in an episode does not interfere with other episodes. Realistic scenarios are nonepisodic, so the robot actually has an advantage in planning and keeping a memory of the past, in order to have an acceptable performance, in a wide range of environments.
- **Static/dynamic**—static environments cannot change while the robot deliberates. Dynamic environments can. Again, robots usually have to face dynamic environments.
- **Discrete/continuous**—Discrete environments have only a limited number of perceptions and actions. Those are, obviously, very rare. Robots in the real world (not simulated) have to deal with continuous variables.

According to these properties, the problem of robot navigation is the worst possible: the environment is not accessible, nondeterministic, nonepisodic, dynamic and continuous. Of course, in simulations or laboratory conditions, some or all of these properties can be controlled. In some particular applications, such as industrial robots in factory environments, the range of expected scenarios can also be limited. But, in general, the environments faced by real robots are the worst possible, and the robots must be robust enough to behave well under those conditions.

### 1.3. Summary

The dream of intelligent machines is thousands of years old: it dates back to the ancient Egypt and to the Greek mythology, at least. The knowledge to build those intelligent ma-



chines comes from old and recent scientific discoveries. The 20<sup>th</sup> century has been specially prolific. The problem is naturally very complex, because the goal is to build machines which have to make successful decisions in very challenging environments. But it is even more difficult due to the fact that the goal is not clear, because the main characteristic that is sought to those machines—intelligence—is still poorly defined.

This book reviews important scientific work about Intelligence, Associative Memories and Robot Navigation based on Visual Memories. It extensively describes an approach for robot navigation which relies on the use of a Sparse Distributed Memory. The robot stores memories of places it visits and it is later able to visit them again by following the same paths, or inferring new routes by detecting connection points between paths. It discards memories which are, with high probability, not useful, and detects connection points and common segments between paths. The limits of the system under changing and unfavourable conditions are also studied and reported.