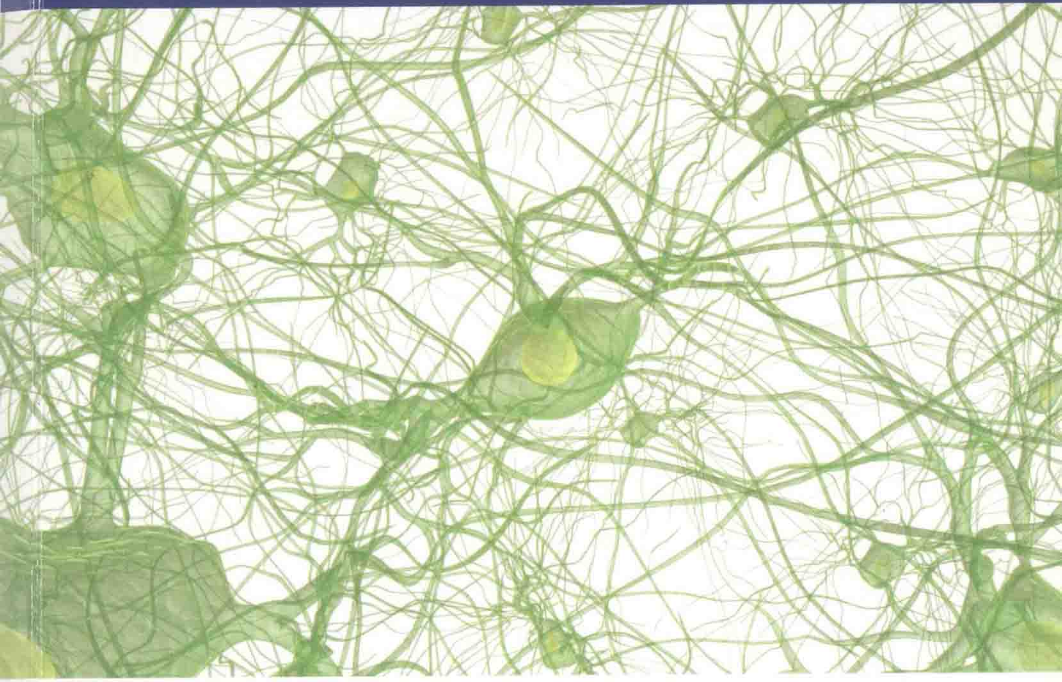

Tyler W. Garaas

Real-Time Low-Level Active Robotic Vision

Based on Biologically Inspired Neural Models

LAP LAMBERT Academic Publishing

Tyler W. Garaas

Real-Time Low-Level Active Robotic Vision

Based on Biologically Inspired Neural Models

LAP LAMBERT Academic Publishing

Impressum/Imprint (nur für Deutschland/ only for Germany)
Bibliografische Information der Deutschen Nationalbibliothek: Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über http://dnb.d-nb.de abrufbar.

Coverbild: www.ingimage.com

Verlag: LAP LAMBERT Academic Publishing GmbH & Co. KG
Dudweiler Landstr. 99, 66123 Saarbrücken, Deutschland
Telefon +49 681 3720-310, Telefax +49 681 3720-3109
Email: info@lap-publishing.com

Herstellung in Deutschland:
Schaltungsdienst Lange o.H.G., Berlin
Books on Demand GmbH, Norderstedt
Reha GmbH, Saarbrücken
Amazon Distribution GmbH, Leipzig
ISBN: 978-3-8433-7630-3

Imprint (only for USA, GB)
Bibliographic information published by the Deutsche Nationalbibliothek: The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Cover image: www.ingimage.com

Publisher: LAP LAMBERT Academic Publishing GmbH & Co. KG
Dudweiler Landstr. 99, 66123 Saarbrücken, Germany
Phone +49 681 3720-310, Fax +49 681 3720-3109
Email: info@lap-publishing.com

Printed in the U.S.A.
Printed in the U.K. by (see last page)
ISBN: 978-3-8433-7630-3

Tyler W. Garaas

Real-Time Low-Level Active Robotic Vision

To my parents, who allowed me the freedom to explore,

sometimes at the expense of a [now disassembled] lawnmower engine or computer.

To my grandparents, who never stopped believing in me,

despite my many attempts to dissuade.

To my wife, who has always supported me,

and on occasion joined me in the sleepless nights.

Acknowledgments

Many people have contributed to the success of the current work. Most prominent was the continuous guidance provided by my advisor, lab director, and friend, Dr. Marc Pomplun. Upon arriving at UMass Boston for my graduate work, I knew little of vision and eye movements; of the latter, I was barely aware of my own. Throughout my time under his guidance, I have learned so much about [robotic and biological] vision, neuroscience, and cognitive science. I look forward to our future collaboration.

I would also like to thank everyone I have worked with in the Visual Attention Laboratory; it was my great pleasure to work with you throughout the years. I'm sure I will never forget our always-stimulating lab outings.

I am also indebted to the support and opportunities provided to me by the Department of Computer Science at the University of Massachusetts Boston. I have taken many ventures throughout my time in the department, and I was never more than a walk down the hall away from assistance in my journey. Notably, I would like to express my gratitude for the support provided by my Graduate Assistance in the Areas of National Need (GAANN) grant, which allowed me to carry out my academic pursuits with relative ease. I hope I have met – and will continue to meet – the goals of the grant.

Finally, I would like to express my gratitude to my dissertation committee. I have not been the easiest of candidates to oversee, but they have continually applied great patients with my many untimely requests. In particular, I am thankful Dr. Paul Nestor could put aside his fear of computers to join my dissertation committee.

Foreword

Complex systems exist in every microscopic corner and every macroscopic plane of the physical world in which we inhabit. Whether it be the transformation of work-energy into heat-energy, decades-long weather patterns, nuclear fusion taking place in the stars, or the emergence of self-organizing life, complex systems remain to us a thing to be understood; a puzzle begging to be solved by clever manipulations of electrochemical signals. Perhaps some of the greatest questions that have been asked are concerned with topics of self-image – where did we come from, how did we get here, and why are we here? Each of these questions remains glaringly unsolved; puzzles that we have made great progress towards solving, but to which no solution has been reached. I must acknowledge the fact that my research – and perhaps the majority of all science – is, on some level, an attempt to make an iota of progress so that one day we might understand.

Contents

List of Figures

List of Tables

List of Code Listings

1 Introduction

1.1 Motivation

At present, there is a high demand for entities that can perform certain tasks in place of a human peer. Some of the most obvious examples have already been implemented; astronauts, for example, no longer make missions to the moon or Mars, but, instead, mechanical entities – though not technically robots, since they are operated by a human agent, I use the term here for convenience – are used to explore the harsh landscapes of such faraway places. This situation makes the tradeoff of supporting the survival of a human explorer for the loss of efficiency that is afforded by a human explorer. Robots are also becoming commonplace in search-and-rescue operations where human rescuers would like to be able to search a volatile area for survivors before risking their lives; such robots were used to search for survivors in the tragic attacks on September 11th; again, a tradeoff between safety and efficiency is prominent. A wealth of examples can immediately be thought of that could also benefit from such a tradeoff. Indeed, robots have already saved the lives of hundreds of humans by diffusing improvised explosive devices in combat situations.

In each of the cases above, the mechanical entity that is used requires a human operator to direct the movements of the entity. If the human operator could be removed from the equation, the amount of work that could be performed by the mechanical entity could increase significantly, as operators require food, sleep, training, and breaks. Indeed, it seems unlikely that a proof of a human's fault tolerance under specific conditions is forthcoming anytime soon. Unfortunately, the level of control mechanisms is incapable of supporting the tasks that would be required of the robot. As such, there is a great need for the development of control algorithms to support the work to be performed by the robot.

Still, robots captured our imagination well over fifty years ago in science fiction novels and movies. These animated companions have been our best friend, our greatest enemy, or, in some cases, our significant other. The robots in such stories often came equipped with super-human capabilities and intelligences which allowed them to save or enslave us mere humans. Opposite these fantastic characters of our imagination are the current generation of robots which instill a much more sober perspective.

One of the components an entity must possess in order to be considered robotic is an ability to sense the outside world; thus, a player piano cannot be considered robotic even though it is an autonomous agent. However, sensory systems can be very complex; at present, the most successful commercial robot navigates its terrain using a simple bumper mechanism in tandem with virtual infrared walls. Although the robot will do a great job at its task, the level of sophistication of its sensory mechanisms provide an indicator of our current climb to meet the vision laid out by our imagination.

Simple sensory mechanisms akin to those of the robot described above appear in biological creatures as well. Many "simpler" creatures rely on tactile senses to tell them when they encounter an object. For instance, jellyfish have well-developed tactile senses that can let them know when to defend, flee, or attack. Insects, in contrast, have a sense of vision, but their ability to see is much closer to the way in which a robot 'sees' the virtual infrared walls than it is to the manner in which we see something. Interestingly, the complexity level of sensory mechanisms in biological creatures correlates well with our own subjective view of their evolutionary progress.

Humans are primarily visual animals, with the other four senses left to fill in any gaps that are left behind. It is our incredibly sophisticated vision system that allows us to efficiently navigate the world, to construct and manipulate tools, and to interact with our fellow man. The importance of vision has not been lost on robotics researchers, as it has been the focus of intense study since we first began to imagine the possibilities of robots. However, the complex nature of vision has stifled the progress of realizing even moderately capable robotic vision systems. As an example of the complex situations that must be handled, consider the case where you are sitting in an office that is lit by two primary, over-head sources. Now, if one of the light sources is removed, the only change in your perception of the office is the change in the amount and direction of the light. This is true despite the fact that none of your retinal cells are receiving the same signal as they did before. How then can we have a robot know that the image it sees of a chair now and an image of a chair it sees later under one of an infinite number of combinations of light intensity and direction is the same chair? The answer is that no one currently has a definitive answer.

Even though no one currently knows the answer to the question of how a robot is to distinguish an object under various viewing conditions, many accomplished researchers have made impressive progress in devising algorithms to perform likewise visual tasks. Most of these algorithms involve some sort of statistical calculation on the pixels of the image and then attempt to determine likely properties of the image based on the computed statistics. This has resulted in some impressive computer vision algorithms such as optic flow analysis, edge detection, and image segmentation. However, these capabilities pale in comparison to those employed by even simple biological organisms, especially in the category of robustness.

If we are unable to devise algorithms that are able to compute the image properties required to produce a capable robotic vision system, how are we supposed to create robots of sufficient ability? One potential answer may come from another group of researchers that are just as interested in figuring out how such tasks can be performed as roboticists are; this group of researchers is the vision neuroscientists. Vision neuroscientists also want to know how it is that one can recognize a chair under different viewing conditions; however, instead of using this knowledge to build robotic vision systems, these researchers want to know how our own brain performs complex visual mechanisms such as this. Indeed, visual neuroscience is a very active area of research from which many important principles regarding vision have been uncovered.

In the present work, I propose a robotic vision system built solely using simulated artificial neurons to perform certain visual tasks that would be beneficial to a larger robotic system that incorporated it. The vision system has two overriding properties that should be considered at each stage of the construction; (i) the system should follow as closely as possible neuroanatomical and neurophysiological data; (ii) individual neurons in the system should implement a known functional purpose. An inherent tradeoff will need to be made at each level of the system, as artificial neurons that very accurately simulate their biological counterpart