NEURAL NETWORK PERCEPTION FOR MOBILE ROBOT GUIDANCE

THE KLUWER INTERNATIONAL SERIES IN ENGINEERING AND COMPUTER SCIENCE

ROBOTICS: VISION, MANIPULATION AND SENSORS

Consulting Editor: **Takeo Kanade**

VISION AND NAVIGATION: THE CMU NAVLAB, C. Thorpe (editor)
ISBN 0-7923-9068-7
TASK-DIRECTED SENSOR FUSION AND PLANNING: A Computational Approach, G. D. Hager
ISBN: 0-7923-9108-X
COMPUTER ANALYSIS OF VISUAL TEXTURES, F. Tomita and S. Tsuji
ISBN: 0-7923-9114-4
DATA FUSION FOR SENSORY INFORMATION PROCESSING SYSTEMS, J. Clark and A. Yuille
ISBN: 0-7923-9120-9
PARALLEL ARCHITECTURES AND PARALLEL ALGORITHMS FOR INTEGRATED VISION SYSTEMS, A.N. Choudhary, J. H. Patel
ISBN: 0-7923-9078-4
ROBOT MOTION PLANNING, J. Latombe
ISBN: 0-7923-9129-2
DYNAMIC ANALYSIS OF ROBOT MANIPULATORS: A Cartesian Tensor Approach, C.A Balafoutis, R.V. Patel
ISBN: 07923-9145-4
PERTURBATION TECHNIQUES FOR FLEXIBLE MANIPULATORS: A. Fraser and R. W. Daniel
ISBN: 0-7923-9162-4
COMPUTER AIDED MECHANICAL ASSEMBLY PLANNING: L. Homen de Mello and S. Lee
ISBN: 0-7923-9205-1
INTELLIGENT ROBOTIC SYSTEMS FOR SPACE EXPLORATION: Alan A. Desrochers
ISBN: 0-7923-9197-7
MEASUREMENT OF IMAGE VELOCITY: David J. Fleet
ISBN: 0-7923-9198-5
DIRECTED SONAR SENSING FOR MOBILE ROBOT NAVIGATION: John J. Leonard, Hugh F. Durrant-Whyte
ISBN: 0-7923-9242-6
A GENERAL MODEL OF LEGGED LOCOMOTION ON NATURAL TERRAIN: David J. Manko
ISBN: 0-7923-9247-7
INTELLIGENT ROBOTIC SYSTEMS: THEORY, DESIGN AND APPLICATIONS K. Valavanis, G. Saridis
ISBN : 0-7923-9250-7
QUALITATIVE MOTION UNDERSTANDING: W. Burger, B. Bhanu
ISBN: 0-7923-9251-5
DIRECTED SONAR SENSING FOR MOBILE ROBOT NAVIGATION: J.J. Leonard, H.F. Durrant-Whyte
ISBN: 0-7923-9242-6
NONHOLONOMIC MOTION PLANNING: Z. Li, J.F. Canny
ISBN: 0-7923-9275-2
SPACE ROBOTICS: DYNAMICS AN D CONTROL: Y. Xu, T. Kanade
ISBN: 0-7923-9265-5
NEURAL NETWORKS IN ROBOTICS: G.A. Bekey, K.Y. Goldberg
ISBN: 0-7923-9268-X
EFFICIENT DYNAMIC SIMULATON OF ROBOT MECHANISMS: K. Lilly
ISBN: 0-7923-9286-8

NEURAL NETWORK PERCEPTION FOR MOBILE ROBOT GUIDANCE

by

Dean A. Pomerleau
Carnegie Mellon University

Springer Science+Business Media, LLC

Library of Congress Cataloging-in-Publication Data

Pomerleau, Dean A., 1964-
Neural network perception for mobile robot guidance / by Dean A. Pomerleau.
p. cm. -- (The Kluwer international series in engineering and computer science ; SECS 239)
Includes bibliographical references and index.
ISBN 978-1-4613-6400-9 ISBN 978-1-4615-3192-0 (eBook)
DOI 10.1007/978-1-4615-3192-0
1. Mobile robots. 2. Neural networks (Computer science)
3. Robots--Control systems. I. Title. II. Series.
TJ211.415.P66 1993
629.8'95--dc20 93-24616
CIP

Copyright © 1993 by Springer Science+Business Media New York
Originally published by Kluwer Academic Publishers in 1993
Softcover reprint of the hardcover 1st edition 1993

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, mechanical, photo-copying, recording, or otherwise, without the prior written permission of the publisher, Springer Science+Business Media, LLC.

Printed on acid-free paper.

Dedicated to Terry, Glen and Phyllis

Contents

Foreword

Dean Pomerleau's trainable road tracker, ALVINN, is arguably the world's most famous neural net application. It currently holds the world's record for distance traveled by an autonomous robot without interruption: 21.2 miles along a highway, in traffic, at speeds of up to 55 miles per hour. Pomerleau's work has received worldwide attention, including articles in Business Week (March 2, 1992), Discover (July, 1992), and German and Japanese science magazines. It has been featured in two PBS series, "The Machine That Changed the World" and "By the Year 2000," and appeared in news segments on CNN, the Canadian news and entertainment program "Live It Up", and the Danish science program "Chaos".

What makes ALVINN especially appealing is that it does not merely drive – it *learns* to drive, by watching a human driver for roughly five minutes. The training inputs to the neural network are a video image of the road ahead and the current position of the steering wheel. ALVINN has learned to drive on single lane, multi-lane, and unpaved roads. It rapidly adapts to other sensors: it learned to drive at night using laser reflectance imaging, and by using a laser rangefinder it learned to swerve to avoid obstacles and maintain a fixed distance from a row of parked cars. It has even learned to drive backwards.

The architecture has been tested on three different platforms. In addition to driving the CMU NAVLAB (a modified Chevy van) and NAVLAB II (a modified HMMWV jeep), ALVINN was used to control SM^2, the Self-Mobile Space Manipulator – a visually-guided robot arm under development for NASA. The SM^2 is designed to walk in zero gee along the trusswork of Space Station Freedom.

It would be a mistake to view ALVINN as just a flashy demo. Pomerleau's dissertation contains substantial technical contributions that significantly advance the state of the art. Specifically, he presents:

- New methodology for "on the fly" training of neural networks. The training set must be kept small to meet real-time processing constraints, but examples should not be discarded until they have been adequately learned. The training set must also be shaped to represent a balanced variety of cases (i.e., equal numbers of left and right turns at various angles), even though the training data the robot actually observes in 5 minutes is badly skewed.

- New methodology for estimating the accuracy of a neural network's response, and for arbitrating between multiple expert networks. The OARE (Output Appearance Reliability Estimation) and IRRE (Input Reconstruction Reliability Estimation) methods solve this vitally important problem, and allow the robot to transition seamlessly between different road types.

- Advances in network analysis and training techniques. Pomerleau uses sensitivity analysis and inspection of weight matrices to derive high-level explanations of the feature detectors the network develops, and the way these detectors interact to collectively produce correct driving behavior. He uses this knowledge to develop an interesting new training technique, structured noise, that can be used to shape a training set to quickly teach the network how to behave better

This thesis is a tour de force, a combination of important theoretical advances and dramatic practical demonstrations. Dean Pomerleau has done more than any other person to establish the credibility of neural networks for real-time learning and visually-guided control.

David S. Touretzky
Senior Research Scientist
Carnegie Mellon University

Preface

Vision based mobile robot guidance has proven difficult for classical machine vision methods because of the diversity and real time constraints inherent in the task. This book describes a connectionist system called ALVINN (Autonomous Land Vehicle In a Neural Network) that overcomes these difficulties. ALVINN learns to guide mobile robots using the back-propagation training algorithm. Because of its ability to learn from example, ALVINN can adapt to new situations and therefore cope with the diversity of the autonomous navigation task.

But real world problems like vision based mobile robot guidance presents a different set of challenges for the connectionist paradigm. Among them are:

- How to develop a general representation from a limited amount of real training data,
- How to understand the internal representations developed by artificial neural networks,
- How to estimate the reliability of individual networks,
- How to combine multiple networks trained for different situations into a single system,
- How to combine connectionist perception with symbolic reasoning.

This book presents novel solutions to each of these problems. Using these techniques, the ALVINN system can learn to control an autonomous van in under 5 minutes by watching a person drive. Once trained, individual ALVINN networks can drive in a variety of circumstances, including single-lane paved and unpaved roads, and multi-lane lined and unlined roads, at speeds of up to 55 miles per hour. The techniques also are shown to generalize to the task of controlling the precise foot placement of a walking robot.

This book is a revised version of my PhD. thesis, submitted in February 1992 to the Carnegie Mellon University School of Computer Science.

I wish to thank my advisor, Dr. David Touretzky for his support and technical advice during my graduate studies. Dave has not only provided invaluable feedback, he has also given me the latitude I've needed to develop as a researcher. I am also grateful to Dr. Charles Thorpe for the opportunities, resources and expertise he has provided me. Without Chuck's support, this research would not have been possible. I also wish to thank the other members of my committee, Dr. Takeo Kanade and Dr. Terrence Sejnowski, for their insightful analysis and valuable comments concerning my work.

I owe much to all the members of the ALV/UGV project. Their technical support and companionship throughout the development of the ALVINN system has made my work both possible and enjoyable. I would like to specifically thank Jay Gowdy and Omead Amidi, whose support software underlies much of the ALVINN system. James

Frazier also deserves thanks for his patience during many hours of test runs on the Navlab.

Interaction with members of the Boltzmann group at Carnegie Mellon has also been indispensable. From them I have not only learned about all aspects of connectionism, but also how to communicate my thoughts and ideas. In particular, the insights and feedback provided by discussions with John Hampshire form the basis for much of this work. I am grateful to Dave Plaut, whose helpful suggestions in the early stages of this work put me on the right track.

Other people who have contributed to the success of this work are the members of the SM^2 group. In particular, Ben Brown and Hiroshi Ueno have given me the opportunity, incentive and support I have needed to explore an alternative domain for connectionist mobile robot guidance.

I am also in debt to my office mates, Spiro Michaylov and Nevin Heintze. They have helped me throughout our time as graduate students, with everything from explaining LaTeX peculiarities to feeding my fish. I would like to thank my parents, Glen and Phyllis, for encouraging my pursuit of higher education, and for all the sacrifices they've made to provide me with a world of opportunities. Finally, I am especially grateful to my wife Terry for her constant love, support and patience during the difficult months spent preparing this book. Her presence in my life has helped me keep everything in perspective.

Dean A. Pomerleau
Carnegie Mellon University

NEURAL NETWORK PERCEPTION FOR MOBILE ROBOT GUIDANCE