
Springer Tracts in Advanced Robotics

Volume 118

Editors: Bruno Siciliano · Oussama Khatib

Peter Corke

Robotics, Vision and Control

Fundamental Algorithms in MATLAB®

Second, completely revised, extended and updated edition

With 492 Images

Additional material is provided at www.petercorke.com/RVC

Professor Bruno Siciliano

Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione,
Università di Napoli Federico II, Via Claudio 21, 80125 Napoli, Italy,
e-mail: siciliano@unina.it

Professor Oussama Khatib

Artificial Intelligence Laboratory, Department of Computer Science,
Stanford University, Stanford, CA 94305-9010, USA,
e-mail: khatib@cs.stanford.edu

Author**Peter Corke**

School of Electrical Engineering and Computer Science
Queensland University of Technology (QUT), Brisbane QLD 4000, Australia
e-mail: rvc@petercorke.com

ISSN 1610-7438
Springer Tracts in Advanced Robotics

ISSN 1610-742X (electronic)

ISBN 978-3-319-54412-0
DOI 10.1007/978-3-319-54413-7

ISBN 978-3-319-54413-7 (eBook)

Library of Congress Control Number: 2017934638

1st ed. 2011 © Springer-Verlag Berlin Heidelberg 2011
© Springer International Publishing AG 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Production: Armin Stasch and Scientific Publishing Services Pvt. Ltd. Chennai, India
Typesetting and layout: Stasch · Bayreuth (stasch@stasch.com)

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Editorial Advisory Board

Nancy Amato, Texas A & M, USA
Oliver Brock, TU Berlin, Germany
Herman Bruyninckx, KU Leuven, Belgium
Wolfram Burgard, Univ. Freiburg, Germany
Raja Chatila, ISIR – UPMC & CNRS, France
Francois Chaumette, INRIA Rennes – Bretagne Atlantique, France
Wan Kyun Chung, POSTECH, Korea
Peter Corke, Queensland Univ. Technology, Australia
Paolo Dario, Scuola S. Anna Pisa, Italy
Alessandro De Luca, Sapienza Univ. Roma, Italy
Rüdiger Dillmann, Univ. Karlsruhe, Germany
Ken Goldberg, UC Berkeley, USA
John Hollerbach, Univ. Utah, USA
Lydia Kavraki, Rice Univ., USA
Vijay Kumar, Univ. Pennsylvania, USA
Bradley Nelson, ETH Zürich, Switzerland
Frank Park, Seoul National Univ., Korea
Tim Salcudean, Univ. British Columbia, Canada
Roland Siegwart, ETH Zurich, Switzerland
Gaurav Sukhatme, Univ. Southern California, USA

More information about this series at <http://www.springer.com/series/5208>

*To my family Phillipa, Lucy and Madeline for their indulgence and support;
my parents Margaret and David for kindling my curiosity;
and to Lou Paul who planted the seed that became this book.*

Foreword

Once upon a time, a very thick document of a dissertation from a faraway land came to me for evaluation. *Visual robot control* was the thesis theme and *Peter Corke* was its author. Here, I am reminded of an excerpt of my comments, which reads, *this is a masterful document, a quality of thesis one would like all of one's students to strive for, knowing very few could attain – very well considered and executed.*

The connection between robotics and vision has been, for over two decades, the central thread of Peter Corke's productive investigations and successful developments and implementations. This rare experience is bearing fruit in this second edition of his book on *Robotics, Vision, and Control*. In its melding of theory and application, this second edition has considerably benefited from the author's unique mix of academic and real-world application influences through his many years of work in robotic mining, flying, underwater, and field robotics.

There have been numerous textbooks in robotics and vision, but few have reached the level of integration, analysis, dissection, and practical illustrations evidenced in this book. The discussion is thorough, the narrative is remarkably informative and accessible, and the overall impression is of a significant contribution for researchers and future investigators in our field. Most every element that could be considered as relevant to the task seems to have been analyzed and incorporated, and the effective use of Toolbox software echoes this thoroughness.

The reader is taken on a realistic walkthrough the fundamentals of mobile robots, navigation, localization, manipulator-arm kinematics, dynamics, and joint-level control, as well as camera modeling, image processing, feature extraction, and multi-view geometry. These areas are finally brought together through extensive discussion of visual servo system. In the process, the author provides insights into how complex problems can be decomposed and solved using powerful numerical tools and effective software.

The *Springer Tracts in Advanced Robotics (STAR)* is devoted to bringing to the research community the latest advances in the robotics field on the basis of their significance and quality. Through a wide and timely dissemination of critical research developments in robotics, our objective with this series is to promote more exchanges and collaborations among the researchers in the community and contribute to further advancements in this rapidly growing field.

Peter Corke brings a great addition to our STAR series with an authoritative book, reaching across fields, thoughtfully conceived and brilliantly accomplished.

Oussama Khatib
Stanford, California
October 2016

Preface

*Tell me and I will forget.
Show me and I will remember.
Involve me and I will understand.*
Chinese proverb

*Simple things should be simple,
complex things should be possible.*
Alan Kay

These are exciting times for robotics. Since the first edition of this book was published we have seen much progress: the rise of the self-driving car, the Mars science laboratory rover making profound discoveries on Mars, the Philae comet landing attempt, and the DARPA Robotics Challenge. We have witnessed the drone revolution – flying machines that were once the domain of the aerospace giants can now be bought for just tens of dollars. All this has been powered by the continuous and relentless improvement in computer power and tremendous advances in low-cost inertial sensors and cameras – driven largely by consumer demand for better mobile phones and gaming experiences. It's getting easier for individuals to create robots – 3D printing is now very affordable, the Robot Operating System (ROS) is both capable and widely used, and powerful hobby technologies such as the Arduino, Raspberry Pi, Dynamixel servo motors and Lego's EV3 brick are available at low cost. This in turn has contributed to the rapid growth of the global maker community – ordinary people creating at home what would once have been done by a major corporation. We have also witnessed an explosion of commercial interest in robotics and computer vision – many startups and a lot of acquisitions by big players in the field. Robotics even featured on the front cover of the Economist magazine in 2014!

So how does a robot work? Robots are data-driven machines. They acquire data, process it and take action based on it. The data comes from sensors measuring the velocity of a wheel, the angle of a robot arm's joint or the intensities of millions of pixels that comprise an image of the world that the robot is observing. For many robotic applications the amount of data that needs to be processed, in real-time, is massive. For a vision sensor it can be of the order of tens to hundreds of megabytes per second.

Progress in robots and machine vision has been, and continues to be, driven by more effective ways to process data. This is achieved through new and more efficient algorithms, and the dramatic increase in computational power that follows Moore's law. ◀ When I started in robotics and vision in the mid 1980s, see Fig. 0.1, the IBM PC had been recently released – it had a 4.77 MHz 16-bit microprocessor and 16 kbytes (expandable to 256 k) of memory. Over the intervening 30 years computing power has perhaps doubled 20 times which is an increase by a factor of one million.

Over the fairly recent history of robotics and machine vision a very large body of algorithms has been developed to efficiently solve large-scale problems in perception, planning, control and localization – a significant, tangible, and collective achievement of the research community. However its sheer size and complexity presents a very real barrier to somebody new entering the field. Given so many algorithms from which to choose, a real and important question is:

What is the right algorithm for this particular problem?

One strategy would be to try a few different algorithms and see which works best for the problem at hand, but this is not trivial and leads to the next question:

How can I evaluate algorithm X on my own data without spending days coding and debugging it from the original research papers?

"Computers in the future may weigh no more than 1.5 tons." Popular Mechanics, forecasting the relentless march of science, 1949




Fig. 0.1.

Once upon a time a lot of equipment was needed to do vision-based robot control. The author with a large rack full of real-time image processing and robot control equipment (1992)

Two developments come to our aid. The first is the availability of general purpose mathematical software which it makes it easy to prototype algorithms. There are commercial packages such as MATLAB®, Mathematica®, Maple® and MathCad®, as well as open source projects include SciLab, Octave, and PyLab. All these tools deal naturally and effortlessly with vectors and matrices, can create complex and beautiful graphics, and can be used interactively or as a programming environment. The second is the open-source movement. Many algorithms developed by researchers are available in open-source form. They might be coded in one of the general purpose mathematical languages just mentioned, or written in a mainstream language like C, C++, Java or Python.

For more than twenty years I have been part of the open-source community and maintained two open-source MATLAB Toolboxes: one for robotics and one for machine vision. They date back to my own Ph.D. work and have evolved since then, growing features and tracking changes to the MATLAB language. The Robotics Toolbox has also been translated into a number of different languages such as Python, SciLab and LabView. More recently some of its functionality is finding its way into the MATLAB Robotics System Toolbox™ published by The MathWorks.

These Toolboxes have some important virtues. Firstly, they have been around for a long time and used by many people for many different problems so the code can be accorded some level of trust. New algorithms, or even the same algorithms coded in new languages or executing in new environments, can be compared against implementations in the Toolbox.

» allow the user to work with real problems, not just trivial examples

Secondly, they allow the user to work with real problems, not just trivial examples. For real robots, those with more than two links, or real images with millions of pixels the computation required is beyond unaided human ability. Thirdly, they allow us to gain insight which can otherwise get lost in the complexity. We can rapidly and easily experiment, play *what if* games, and depict the results graphically using the powerful 2D and 3D graphical display tools of MATLAB. Fourthly, the Toolbox code makes many common algorithms tangible and accessible. You can read the code, you can apply it to your own problems, and you can extend it or rewrite it. It gives you a “leg up” as you begin your journey into robotics.

» a narrative that covers robotics and computer vision
 – both separately and together

Respectively the trademarks of The MathWorks Inc., Wolfram Research, MapleSoft and PTC.

The term machine vision is uncommon today, but it implied the use of real-time computer vision techniques in an industrial setting for some monitoring or control purpose. For robotics the real-time aspect is critical but today the interesting challenges are in nonindustrial applications such as outdoor robotics. The term robotic vision is gaining currency and is perhaps a modern take on machine vision.

The genesis of the book lies in the tutorials and reference material that originally shipped with the Toolboxes from the early 1990s, and a conference paper describing the Robotics Toolbox that was published in 1995. After a false start in 2004, the first edition of this book was written in 2009–2010. The book takes a conversational approach, weaving text, mathematics and examples into a narrative that covers robotics and computer vision – both separately and together. I wanted to show how complex problems can be decomposed and solved using just a few simple lines of code. More formally this is an inductive learning approach, going from specific and concrete examples to the more general.

» show how complex problems can be decomposed and solved

The topics covered in this book are based on my own interests but also guided by real problems that I observed over many years as a practitioner of both robotics and computer vision. I want to give the reader a flavor of what robotics and vision is about and what it can do – consider it a grand tasting menu. I hope that by the end of this book you will share my enthusiasm for these topics.

» consider it a grand tasting menu

I was particularly motivated to present a solid introduction to computer vision for roboticists. The treatment of vision in robotics textbooks tends to concentrate on simple binary vision techniques. In this book we will cover a broad range of topics including color vision, advanced segmentation techniques, image warping, stereo vision, motion estimation, bundle adjustment, visual odometry and image retrieval. We also cover nonperspective imaging using fisheye lenses, catadioptric optics and the emerging area of light-field cameras. These topics are growing in importance for robotics but are not commonly covered. Vision is a powerful sensor, and roboticists should have a solid grounding in modern fundamentals. The last part of the book shows how vision can be used as the primary sensor for robot control.

This book is unlike other text books, and deliberately so. Firstly, there are already a number of excellent text books that cover robotics and computer vision separately and in depth, but few that cover both in an integrated fashion. Achieving such integration is a principal goal of the book.

» software is a first-class citizen in this book

Secondly, software is a first-class citizen in this book. Software is a tangible instantiation of the algorithms described – it can be read and it can be pulled apart, modified and put back together again. There are a number of classic books that use software in an illustrative fashion and have influenced my approach, for example *LaTeX: A document preparation system* (Lamport 1994), *Numerical Recipes in C* (Press et al. 2007), *The Little Lisper* (Friedman et al. 1987) and *Structure and Interpretation of Classical Mechanics* (Sussman et al. 2001). Over 1 000 examples in this book illustrate how the Toolbox software can be used and generally provide *instant gratification* in just a couple of lines of MATLAB code.

» instant gratification in just a couple of lines of MATLAB code

Thirdly, building the book around MATLAB and the Toolboxes means that we are able to tackle more realistic and more complex problems than other books.

» this book provides a complementary approach

The emphasis on software and examples does not mean that rigor and theory are unimportant – they are very important, but this book provides a complementary approach. It is best read in conjunction with standard texts which do offer rigor and theoretical nourishment. The end of each chapter has a section on further reading and provides pointers to relevant textbooks and key papers. I try hard to use the least amount of mathematical notation required, if you seek deep mathematical rigor this may not be the book for you.

Writing this book provided the impetus to revise and extend the Toolboxes and to include some great open-source software. I am grateful to the following for code that has been either incorporated into the Toolboxes or which has been wrapped into the Toolboxes. Robotics Toolbox contributions include: mobile robot localization and mapping by Paul Newman; a quadrotor simulator by Paul Pounds; a Symbolic Manipulator Toolbox by Jörn Malzahn; pose-graph SLAM code by Giorgio Grisetti and 3D robot models from the ARTE Robotics Toolbox by Arturo Gil. Machine Vision Toolbox contributions include: RANSAC code by Peter Kovesi; pose estimation by Francesco Moreno-Noguer, Vincent Lepetit, and Pascal Fua; color space conversions by Pascal Getreuer; numerical routines for geometric vision by various members of the Visual Geometry Group at Oxford (from the web site of the Hartley and Zisserman book; Hartley and Zisserman 2003); k -means, SIFT and MSER algorithms from the wonderful VLFeat suite (vlfeat.org); graph-based image segmentation software by Pedro Felzenszwalb; and the OpenSURF feature detector by Dirk-Jan Kroon. The Camera Calibration Toolbox by Jean-Yves Bouguet is used unmodified.

Along the way I became fascinated by the mathematicians, scientists and engineers whose work, hundreds of years ago, underpins the science of robotic and computer vision today. Some of their names have become adjectives like Coriolis, Gaussian, Laplacian or Cartesian; nouns like Jacobian, or units like Newton and Coulomb. They are interesting characters from a distant era when science was a hobby and their day jobs were as doctors, alchemists, gamblers, astrologers, philosophers or mercenaries. In order to know whose shoulders we are standing on I have included small vignettes about the lives of some of these people – a smattering of history as a backstory.

In my own career I have had the good fortune to work with many wonderful people who have inspired and guided me. Long ago at the University of Melbourne John Anderson fired my interest in control and Graham Holmes tried with mixed success to have me “think before I code”. Early on I spent a life-direction-changing ten months working with Richard (Lou) Paul in the GRASP laboratory at the University of Pennsylvania in the period 1988–1989. The genesis of the Toolboxes was my Ph.D. research (1991–1994) and my advisors Malcolm Good (University of Melbourne) and Paul Dunn (CSIRO) asked me good questions and guided my research. Laszlo Nemes (CSIRO) provided great wisdom about life and the ways of organizations, and encouraged me to publish and to open-source my software. Much of my career was spent at CSIRO where I had the privilege and opportunity to work on a diverse range of real robotics projects and to work with a truly talented set of colleagues and friends. Part way through writing the first edition I joined the Queensland University of Technology which made time available to complete that work, and in 2015 sabbatical leave to complete the second.

Many people have helped me in my endeavor and I thank them. I was generously hosted for periods of productive writing at Oxford (both editions) by Paul Newman, and at MIT (first edition) by Daniela Rus. Daniela, Paul and Cédric Pradalier made constructive suggestions and comments on early drafts of that edition. For the second edition I was helped by comments on draft chapters by: Tim Barfoot, Dmitry Bratanov, Duncan Campbell, Donald Dansereau, Tom Drummond, Malcolm Good, Peter Kujala, Obadiah Lam, Jörn Malzahn, Felipe Nascimento Martins, Ajay Pandey, Cédric Pradalier, Dan Richards, Daniela Rus, Sareh Shirazi, Surya Singh, Ryan Smith, Ben Talbot, Dorian Tsai and Ben Upcroft; and assisted with wisdom and content by: François Chaumette, Donald Dansereau, Kevin Lynch, Robert Mahony and Frank Park.

I have tried my hardest to eliminate errors but inevitably some will remain. Please email bug reports to me at rvc@petercorke.com as well as suggestions for improvements and extensions.

Writing the second edition was financially supported by EPSRC Platform Grant EP/M019918/1, QUT Science & Engineering Faculty sabbatical grant, QUT Vice Chancellor's Excellence Award, QUT Robotics and Autonomous Systems discipline and the ARC Centre of Excellence for Robotic Vision (grant CE140100016).

Over both editions I have enjoyed wonderful support from MathWorks, through their author program, and from Springer. My editor Thomas Ditzinger has been a great supporter of this project and Armin Stasch, with enormous patience and dedication in layout and typesetting, has transformed my untidy ideas into a thing of beauty.

Finally, my deepest thanks are to Phillipa who has supported me and "the book" with grace and patience for a very long time and in many different places – without her this book could never have been written.

Peter Corke
Brisbane,
Queensland
October 2016

Note on the Second Edition

It seems only yesterday that I turned in the manuscript for the first edition of this book, but it was in fact December 2010, the end of 20 months of writing. So the oldest parts of the book are over 6 years old – it's time for an update!

The revision principle was to keep the good (narrative style, code as a first-class citizen, soft plastic cover) and eliminate the bad (errors and missing topics). I started with the collected errata for the first edition and pencilled markup from a battered copy of the first edition that I've carried around for years. There were more errors than I would have liked and I thank everybody who submitted errata and suggested improvements.

The first edition was written before I taught in the university classroom or created the MOOCs, which is the inverse of the way books are normally developed. Preparing for teaching gave me insights into better ways to present some topics, particularly around pose representation, robot kinematics and dynamics so the presentation has been adjusted accordingly.

New content includes matrix exponential notation; the basics of screw theory and Lie algebra; inertial navigation; differential steer and omnidirectional mobile robots; a deeper treatment of SLAM systems including scan matching and pose graphs; greater use of MATLAB computer algebra; operational space control; deeper treatment of manipulator dynamics and control; visual SLAM and visual odometry; structured light; bundle adjustment; and light-field cameras.

In the first edition I shied away from Lie algebra, matrix exponentials and twists but I think it's important to cover them. The topic is deeply mathematical and I've tried to steer a middle ground between hardcore algebraic topology and the homogeneous transformation only approach of most other texts, while also staying true to the overall approach of this book.

All MATLAB generated figures have been regenerated to reflect recent improvements to MATLAB graphics and all code examples have been updated as required and tested, and are available as MATLAB Live Scripts.

The second edition of the book is matched by new major releases of my Toolboxes: Robotics Toolbox (release 10) and the Machine Vision Toolbox (release 4). These newer versions of the toolboxes have some minor incompatibilities with previous releases of the toolboxes, and therefore also with the code examples in the first edition of the book.

Contents

1	Introduction	1
1.1	Robots, Jobs and Ethics	7
1.2	About the Book	8
1.2.1	MATLAB Software and the Toolboxes	9
1.2.2	Notation, Conventions and Organization	10
1.2.3	Audience and Prerequisites	11
1.2.4	Learning with the Book	11
1.2.5	Teaching with the Book	12
1.2.6	Outline	12
	Further Reading	14
	Part I Foundations	15
2	Representing Position and Orientation	17
2.1	Working in Two Dimensions (2D)	22
2.1.1	Orientation in 2-Dimensions	23
2.1.2	Pose in 2-Dimensions	26
2.2	Working in Three Dimensions (3D)	31
2.2.1	Orientation in 3-Dimensions	32
2.2.2	Pose in 3-Dimensions	46
2.3	Advanced Topics	49
2.3.1	Normalization	49
2.3.2	Understanding the Exponential Mapping	50
2.3.3	More About Twists	52
2.3.4	Dual Quaternions	55
2.3.5	Configuration Space	55
2.4	Using the Toolbox	56
2.5	Wrapping Up	58
	Further Reading	60
	Exercises	61
3	Time and Motion	63
3.1	Time-Varying Pose	63
3.1.1	Derivative of Pose	63
3.1.2	Transforming Spatial Velocities	64
3.1.3	Incremental Rotation	66
3.1.4	Incremental Rigid-Body Motion	67
3.2	Accelerating Bodies and Reference Frames	68
3.2.1	Dynamics of Moving Bodies	68
3.2.2	Transforming Forces and Torques	69
3.2.3	Inertial Reference Frame	69
3.3	Creating Time-Varying Pose	70
3.3.1	Smooth One-Dimensional Trajectories	70

3.3.2	Multi-Dimensional Trajectories	73
3.3.3	Multi-Segment Trajectories	74
3.3.4	Interpolation of Orientation in 3D	75
3.3.5	Cartesian Motion in 3D	77
3.4	Application: Inertial Navigation	79
3.4.1	Gyroscopes	79
3.4.2	Accelerometers	81
3.4.3	Magnetometers	85
3.4.4	Sensor Fusion	87
3.5	Wrapping Up	90
	Further Reading	90
	Exercises	91
	Part II Mobile Robots	93
4	Mobile Robot Vehicles	99
4.1	Wheeled Mobile Robots	99
4.1.1	Car-Like Mobile Robots	99
4.1.2	Differentially-Steered Vehicle	109
4.1.3	Omnidirectional Vehicle	112
4.2	Flying Robots	114
4.3	Advanced Topics	119
4.3.1	Nonholonomic and Under-Actuated Systems	119
4.4	Wrapping Up	121
	Further Reading	122
	Toolbox and MATLAB Notes	123
	Exercises	123
5	Navigation	125
5.1	Reactive Navigation	126
5.1.1	Braitenberg Vehicles	126
5.1.2	Simple Automata	128
5.2	Map-Based Planning	130
5.2.1	Distance Transform	130
5.2.2	D*	134
5.2.3	Introduction to Roadmap Methods	136
5.2.4	Probabilistic Roadmap Method (PRM)	137
5.2.5	Lattice Planner	140
5.2.6	Rapidly-Exploring Random Tree (RRT)	144
5.3	Wrapping Up	146
	Further Reading	147
	Resources	148
	MATLAB Notes	148
	Exercises	148
6	Localization	151
6.1	Dead Reckoning	155
6.1.1	Modeling the Vehicle	155
6.1.2	Estimating Pose	157
6.2	Localizing with a Map	160
6.3	Creating a Map	165
6.4	Localization and Mapping	167
6.5	Rao-Blackwellized SLAM	169
6.6	Pose Graph SLAM	170

6.7	Sequential Monte-Carlo Localization	175
6.8	Application: Scanning Laser Rangefinder	178
	Laser Odometry	179
	Laser-Based Map Building	181
	Laser-Based Localization	182
6.9	Wrapping Up	182
	Further Reading	183
	Toolbox and MATLAB Notes	185
	Exercises	185
	Part III Arm-Type Robots	189
7	Robot Arm Kinematics	193
7.1	Forward Kinematics	193
	7.1.1 2-Dimensional (Planar) Robotic Arms	194
	7.1.2 3-Dimensional Robotic Arms	196
7.2	Inverse Kinematics	205
	7.2.1 2-Dimensional (Planar) Robotic Arms	205
	7.2.2 3-Dimensional Robotic Arms	207
7.3	Trajectories	211
	7.3.1 Joint-Space Motion	211
	7.3.2 Cartesian Motion	214
	7.3.3 Kinematics in Simulink	214
	7.3.4 Motion through a Singularity	215
	7.3.5 Configuration Change	216
7.4	Advanced Topics	217
	7.4.1 Joint Angle Offsets	217
	7.4.2 Determining Denavit-Hartenberg Parameters	217
	7.4.3 Modified Denavit-Hartenberg Parameters	218
7.5	Applications	220
	7.5.1 Writing on a Surface	220
	7.5.2 A Simple Walking Robot	221
7.6	Wrapping Up	225
	Further Reading	226
	MATLAB and Toolbox Notes	227
	Exercises	227
8	Manipulator Velocity	229
8.1	Manipulator Jacobian	229
	8.1.1 Jacobian in the World Coordinate Frame	229
	8.1.2 Jacobian in the End-Effector Coordinate Frame	232
	8.1.3 Analytical Jacobian	232
8.2	Jacobian Condition and Manipulability	234
	8.2.1 Jacobian Singularities	234
	8.2.2 Manipulability	235
8.3	Resolved-Rate Motion Control	237
	8.3.1 Jacobian Singularity	240
8.4	Under- and Over-Actuated Manipulators	240
	8.4.1 Jacobian for Under-Actuated Robot	241
	8.4.2 Jacobian for Over-Actuated Robot	242
8.5	Force Relationships	244
	8.5.1 Transforming Wrenches to Joint Space	244
	8.5.2 Force Ellipsoids	244
8.6	Inverse Kinematics: a General Numerical Approach	245
	8.6.1 Numerical Inverse Kinematics	245

8.7	Advanced Topics	247
8.7.1	Computing the Manipulator Jacobian Using Twists	247
8.8	Wrapping Up	247
	Further Reading	248
	MATLAB and Toolbox Notes	248
	Exercises	248
9	Dynamics and Control	251
9.1	Independent Joint Control	251
9.1.1	Actuators	251
9.1.2	Friction	252
9.1.3	Effect of the Link Mass	253
9.1.4	Gearbox	254
9.1.5	Modeling the Robot Joint	255
9.1.6	Velocity Control Loop	257
9.1.7	Position Control Loop	261
9.1.8	Independent Joint Control Summary	262
9.2	Rigid-Body Equations of Motion	263
9.2.1	Gravity Term	264
9.2.2	Inertia Matrix	266
9.2.3	Coriolis Matrix	267
9.2.4	Friction	268
9.2.5	Effect of Payload	268
9.2.6	Base Force	269
9.2.7	Dynamic Manipulability	269
9.3	Forward Dynamics	271
9.4	Rigid-Body Dynamics Compensation	272
9.4.1	Feedforward Control	273
9.4.2	Computed Torque Control	274
9.4.3	Operational Space Control	275
9.5	Applications	276
9.5.1	Series-Elastic Actuator (SEA)	276
9.6	Wrapping Up	278
	Further Reading	278
	Exercises	280
	Part IV Computer Vision	283
10	Light and Color	287
10.1	Spectral Representation of Light	287
10.1.1	Absorption	289
10.1.2	Reflectance	290
10.1.3	Luminance	290
10.2	Color	291
10.2.1	The Human Eye	292
10.2.2	Measuring Color	294
10.2.3	Reproducing Colors	295
10.2.4	Chromaticity Space	298
10.2.5	Color Names	300
10.2.6	Other Color and Chromaticity Spaces	301
10.2.7	Transforming between Different Primaries	304
10.2.8	What Is White?	306
10.3	Advanced Topics	306
10.3.1	Color Temperature	306
10.3.2	Color Constancy	307

10.3.3	White Balancing	308
10.3.4	Color Change Due to Absorption	308
10.3.5	Dichromatic Reflectance	310
10.3.6	Gamma	310
10.4	Application: Color Image	312
10.4.1	Comparing Color Spaces	312
10.4.2	Shadow Removal	313
10.5	Wrapping Up	315
	Further Reading	316
	Data Sources	316
	Exercises	317
11	Image Formation	319
11.1	Perspective Camera	319
11.1.1	Perspective Projection	319
11.1.2	Modeling a Perspective Camera	322
11.1.3	Discrete Image Plane	324
11.1.4	Camera Matrix	325
11.1.5	Projecting Points	327
11.1.6	Lens Distortion	330
11.2	Camera Calibration	331
11.2.1	Homogeneous Transformation Approach	331
11.2.2	Decomposing the Camera Calibration Matrix	333
11.2.3	Pose Estimation	334
11.2.4	Camera Calibration Toolbox	335
11.3	Wide Field-of-View Imaging	336
11.3.1	Fisheye Lens Camera	337
11.3.2	Catadioptric Camera	340
11.3.3	Spherical Camera	342
11.4	Unified Imaging	344
11.4.1	Mapping Wide-Angle Images to the Sphere	345
11.4.2	Mapping from the Sphere to a Perspective Image	347
11.5	Novel Cameras	348
11.5.1	Multi-Camera Arrays	348
11.5.2	Light-Field Cameras	348
11.6	Advanced Topics	350
11.6.1	Projecting 3D Lines and Quadrics	350
11.6.2	Nonperspective Cameras	352
11.7	Wrapping Up	353
	Further Reading and Resources	354
	Toolbox Notes	355
	Exercises	356
12	Images and Image Processing	359
12.1	Obtaining an Image	359
12.1.1	Images from Files	359
12.1.2	Images from an Attached Camera	363
12.1.3	Images from a Movie File	365
12.1.4	Images from the Web	366
12.1.5	Images from Maps	367
12.1.6	Images from Code	367
12.2	Image Histograms	369
12.3	Monadic Operations	370

12.4	Diadic Operations	372
12.5	Spatial Operations	376
12.5.1	Linear Spatial Filtering	376
12.5.2	Template Matching	387
12.5.3	Nonlinear Operations	392
12.6	Mathematical Morphology	393
12.6.1	Noise Removal	396
12.6.2	Boundary Detection	398
12.6.3	Hit or Miss Transform	398
12.6.4	Distance Transform	399
12.7	Shape Changing	401
12.7.1	Cropping	401
12.7.2	Image Resizing	402
12.7.3	Image Pyramids	403
12.7.4	Image Warping	404
12.8	Wrapping Up	407
	Further Reading	407
	Sources of Image Data	409
	MATLAB Notes	409
	General Software Tools	409
	Exercises	410
13	Image Feature Extraction	413
13.1	Region Features	415
13.1.1	Classification	415
13.1.2	Representation	424
13.1.3	Description	427
13.1.4	Summary	437
13.2	Line Features	438
13.2.1	Summary	443
13.3	Point Features	443
13.3.1	Classical Corner Detectors	443
13.3.2	Scale-Space Corner Detectors	449
13.4	Wrapping Up	454
	MATLAB Notes	454
	Further Reading	455
	Exercises	457
14	Using Multiple Images	459
14.1	Feature Correspondence	460
14.2	Geometry of Multiple Views	464
14.2.1	The Fundamental Matrix	466
14.2.2	The Essential Matrix	468
14.2.3	Estimating the Fundamental Matrix from Real Image Data	470
14.2.4	Planar Homography	474
14.3	Stereo Vision	479
14.3.1	Sparse Stereo	479
14.3.2	Dense Stereo Matching	483
14.3.3	Peak Refinement	489
14.3.4	Cleaning up and Reconstruction	491
14.3.5	3D Texture Mapped Display	494
14.3.6	Anaglyphs	495
14.3.7	Image Rectification	496
14.4	Bundle Adjustment	497

14.5	Point Clouds	503
14.5.1	Fitting a Plane	503
14.5.2	Matching Two Sets of Points	505
14.6	Structured Light	507
14.7	Applications	509
14.7.1	Perspective Correction	509
14.7.2	Mosaicing	512
14.7.3	Image Matching and Retrieval	514
14.7.4	Visual Odometry	520
14.8	Wrapping Up	523
	MATLAB and Toolbox Notes	524
	Further Reading	524
	Resources	528
	Exercises	529
	Part V Robotics, Vision and Control	533
15	Vision-Based Control	537
15.1	Position-Based Visual Servoing	538
15.2	Image-Based Visual Servoing	541
15.2.1	Camera and Image Motion	542
15.2.2	Controlling Feature Motion	547
15.2.3	Estimating Feature Depth	551
15.2.4	Performance Issues	554
15.3	Using Other Image Features	556
15.3.1	Line Features	556
15.3.2	Circle Features	557
15.3.3	Photometric Features	559
15.4	Wrapping Up	560
	Further Reading	560
	Exercises	562
16	Advanced Visual Servoing	565
16.1	XY/Z-Partitioned IBVS	565
16.2	IBVS Using Polar Coordinates	568
16.3	IBVS for a Spherical Camera	570
16.4	Applications	572
16.4.1	Arm-Type Robot	572
16.4.2	Mobile Robot	573
16.4.3	Aerial Robot	576
16.5	Wrapping Up	578
	Further Reading	578
	Resources	579
	Exercises	579
	Appendices	581
A	Installing the Toolboxes	583
B	Linear Algebra Refresher	587
C	Geometry	595
D	Lie Groups and Algebras	611
E	Linearization, Jacobians and Hessians	617
F	Solving Systems of Equations	621
G	Gaussian Random Variables	631
H	Kalman Filter	635
I	Graphs	641
J	Peak Finding	645

Bibliography 649

Index 663

Index of People 663

Index of Functions, Classes and Methods 664

General Index 669

Nomenclature

The notation used in robotics and computer vision varies considerably across books and research papers. The symbols used in this book, and their units where appropriate, are listed below. Some symbols have multiple meanings and their context must be used to disambiguate them.

Notation	Description
x^*	desired value of x
x^+	predicted value of x
$x^\#$	measured, or observed, value of x
\hat{x}	estimated value of x
\bar{x}	mean of x or relative value
$x^{(k)}$	k^{th} element of a time series
\boldsymbol{v}	a vector
$\hat{\boldsymbol{v}}$	a unit-vector parallel to \boldsymbol{v}
$\tilde{\boldsymbol{v}}$	homogeneous representation of vector \boldsymbol{v}
$v[i]$	i^{th} element of vector \boldsymbol{v}
v_x	a component of a vector
\boldsymbol{A}	a matrix
$A[i, j]$	the element (i, j) of \boldsymbol{A}
$A_{i,j}$	the element (i, j) of \boldsymbol{A}
$f(x)$	a function of x
$F_x(x)$	the derivative $\partial f / \partial x$
$F_{xy}(x, y)$	the derivative $\partial^2 f / \partial x \partial y$
\hat{q}	unit quaternion, $\hat{q} \in \mathbb{S}^3$
$\mathbf{0}_{m \times n}$	an $m \times n$ matrix of zeros
$\mathbf{1}_{m \times n}$	an $m \times n$ matrix of ones

Symbol	Description	Unit
B	viscous friction coefficient	N m s rad^{-1}
B	magnetic field intensity (or magnetic flux density)	T
C	camera matrix, $C \in \mathbb{R}^{3 \times 4}$	
$C(\mathbf{q}, \dot{\mathbf{q}})$	manipulator centripetal and Coriolis term	$\text{kg m}^2 \text{s}^{-1}$
\mathcal{C}	configuration space of a robot with N joints: $\mathcal{C} \subset \mathbb{R}^N$	
E	illuminance (lux)	lx
f	focal length	m
f	force	N
\mathbf{f}	vector of image features	
$F(\dot{\mathbf{q}})$	friction torque	Nm
$G(\mathbf{q})$	manipulator gravity loading term	Nm
\mathbb{H}	the set of all quaternions (H for Hamilton)	
$I_{n \times n}$	$n \times n$ identity matrix	
J	inertia	kg m^2
J	inertia tensor, $J \in \mathbb{R}^{3 \times 3}$	kg m^2
J	Jacobian matrix	
${}^A J_B$	Jacobian transforming velocities in frame B to frame A	
k, K	constant	
K	camera calibration matrix	
K_i	amplifier gain (transconductance)	A V^{-1}
K_m	motor torque constant	Nm A^{-1}
L	luminance (nit)	nt
m	mass	kg
$M(\mathbf{q})$	manipulator inertia matrix	kg m^2
$N(\mu, \sigma^2)$	a normal (Gaussian) distribution with mean μ and standard deviation σ	
\mathbf{p}	an image plane point, $\mathbf{p} \in \mathbb{R}^2$	
\mathbf{P}	a world point, $\mathbf{P} \in \mathbb{R}^3$	
\mathbb{P}^2	projective space of all 2-D points, a 3-tuple	
\mathbb{P}^3	projective space of all 3-D points, a 4-tuple	
\mathbf{q}	generalized coordinates, configuration $\mathbf{q} \in \mathcal{C}$	m, rad
\mathbf{Q}	generalized force $\mathbf{Q} \in \mathbb{R}^N$	N, Nm
R	an orthonormal rotation matrix, $R \in \text{SO}(2)$ or $\text{SO}(3)$	
\mathbb{R}	set of real numbers	
\mathbb{R}^2	set of all 2-D points	
\mathbb{R}^3	set of all 3-D points	
s	Laplace transform operator	
\mathbb{S}^1	unit circle, set of angles $[0, 2\pi)$	
\mathbb{S}^n	unit sphere embedded in \mathbb{R}^{n+1}	
$\mathfrak{se}(n)$	Lie algebra for $\text{SE}(n)$, an $\mathbb{R}^{(n+1) \times (n+1)}$ augmented skew-symmetric matrix	
$\mathfrak{so}(n)$	Lie algebra for $\text{SO}(n)$, an $\mathbb{R}^{n \times n}$ skew-symmetric matrix	
$\text{SE}(n)$	special Euclidean group, the set of all poses in n dimensions, represented by an $\mathbb{R}^{(n+1) \times (n+1)}$ homogeneous transformation matrix	
$\text{SO}(n)$	special orthogonal group, the set of all orientations in n dimensions, represented by an $\mathbb{R}^{n \times n}$ orthogonal matrix	
S	twist in 3 dimensions, $S \in \mathbb{R}^6$	
t	time	s
\mathcal{T}	task space of robot: $\mathcal{T} \subset \text{SE}(3)$	K

Symbol	Description	Unit
T	sample interval	s
T	temperature	K
T	optical transmission	m^{-1}
T	homogeneous transformation, $T \in \text{SE}(2)$ or $\text{SE}(3)$	
${}^A T_B$	homogeneous transform representing frame $\{B\}$ with respect to frame $\{A\}$. If A is not given then assumed relative to world coordinate frame 0. Note that ${}^A T_B = ({}^B T_A)^{-1}$	
u, v	camera image plane coordinates	pixels
u_0, v_0	coordinates of the principal point	pixels
\bar{u}, \bar{v}	normalized image plane coordinates, relative to the principal point	m
v	velocity	m s^{-1}
\mathbf{v}	velocity vector	m s^{-1}
\mathbf{W}	wrench, a vector of forces and moments $(f_x, f_y, f_z, m_x, m_y, m_z)$	N, Nm
X, Y, Z	Cartesian coordinates	
\bar{x}, \bar{y}	normalized image-plane coordinates	
\mathbb{Z}	set of all integers	
\mathbb{Z}^+	the set of all integers greater than zero	
ϕ	luminous flux (lumens)	lm
γ	robot steering angle	rad
$\mathbf{\Gamma}$	3-angle representation of rotation, $\mathbf{\Gamma} \in \mathbb{R}^3$	rad
$\mathbf{\Gamma}$	body torque, $\mathbf{\Gamma} \in \mathbb{R}^3$	Nm
θ	angle	rad
$\theta_r, \theta_p, \theta_y$	roll pitch yaw angles	rad
λ	wavelength	m
λ	an eigenvalue	
ν	innovation	
ν	spatial velocity, $\nu = (v_x, v_y, v_z, \omega_x, \omega_y, \omega_z) \in \mathbb{R}^6$	m s^{-1} , rad s^{-1}
ξ	abstract representation of Cartesian pose (pronounced ksi)	
${}^A \xi_B$	abstract representation of relative pose, frame $\{B\}$ with respect to frame $\{A\}$ or rigid-body motion from frame $\{A\}$ to $\{B\}$	
π	mathematic constant	
π	a plane	
ρ_w, ρ_h	pixel width and height	m
σ	standard deviation	
σ	robot joint type, $\sigma = \text{R}$ for revolute and $\sigma = \text{P}$ for prismatic	
Σ	Lie algebra $\Sigma = [\cdot] \in \mathfrak{se}(3)$	
τ	torque	N m
τ_C	Coulomb friction torque	N m
ω	rotational rate	rad s^{-1}
$\boldsymbol{\omega}$	angular velocity vector	rad s^{-1}
ϖ	rotational speed of a motor or propellor	rad s^{-1}
Ω	Lie algebra $\Omega = [\cdot]_{\times} \in \mathfrak{so}(3)$	

Operator	Description	MATLAB
$\ \cdot\ $	norm, or length, of vector: $\mathbb{R}^n \mapsto \mathbb{R}$	<code>norm</code> , <code>.norm</code>
$\mathbf{v}_1 \cdot \mathbf{v}_2$	dot, or inner, product, also $\mathbf{v}_1^T \mathbf{v}_2$: $\mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$	<code>dot</code>
$\mathbf{v}_1 \times \mathbf{v}_2$	cross, or vector, product: $\mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}^n$	<code>cross</code>
\mathbf{A}^{-1}	inverse of \mathbf{A} : $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	<code>inv</code>
\mathbf{A}^+	pseudo-inverse of \mathbf{A} : $\mathbb{R}^{n \times m} \mapsto \mathbb{R}^{m \times n}$	<code>pinv</code>
\mathbf{A}^*	adjugate of $\mathbf{A} \mapsto \det(\mathbf{A})\mathbf{A}^{-1}$, $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	
\mathbf{A}^T	transpose of \mathbf{A} : $\mathbb{R}^{n \times m} \mapsto \mathbb{R}^{m \times n}$	<code>'</code>
\mathbf{A}^{-T}	transpose of inverse $\mathbf{A} \mapsto (\mathbf{A}^T)^{-1} = (\mathbf{A}^{-1})^T$, $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	
\bullet	transform a point (coordinate vector) by a relative pose: $\mathbf{SE}(n) \times \mathbb{R}^n \mapsto \mathbb{R}^n$	<code>*</code>
\oplus	composition: $\mathbf{S}_E^0(n) \times \mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>*</code>
\ominus	composition with inverse: $\mathbf{S}_E^0(n) \times \mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>/</code>
\ominus	unary inverse: $\mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>.inv</code>
$\Delta(\cdot)$	maps incremental pose change to differential motion: $\mathbf{SE}(3) \mapsto \mathbb{R}^6$	<code>tr2delta</code>
$\Delta^{-1}(\cdot)$	maps differential motion to incremental pose change: $\mathbb{R}^6 \mapsto \mathbf{SE}(3)$	<code>delta2tr</code>
$\mathcal{R}_i(\theta)$	pure rotation about axis i : $\mathbb{R} \mapsto \mathbf{SE}(3)$	<code>SE3.rotx y z</code>
$\mathcal{R}(\omega)$	pure rotation by $\ \omega\ $ about ω : $\mathbb{R}^3 \mapsto \mathbf{SE}(3)$	<code>SE3.angvec</code>
$\mathcal{T}_i(d)$	pure translation along axis i : $\mathbb{R} \mapsto \mathbf{SE}(2), \mathbf{SE}(3)$	<code>SE2</code> , <code>SE3</code>
$\mathcal{T}(\mathbf{t})$	pure translation by vector: $\mathbb{R}^n \mapsto \mathbf{SE}(n)$	<code>SE2</code> , <code>SE3</code>
$[\cdot]_t$	translational component of pose: $\mathbf{SE}(n) \mapsto \mathbb{R}^n$	<code>.t</code>
$[\cdot]_R$	rotational component of pose: $\mathbf{SE}(n) \mapsto \mathbb{R}^{n \times n}$	<code>.R</code>
$[\cdot]_\times$	skew-symmetric matrix: $\mathbb{R} \mapsto \mathbf{so}(2), \mathbb{R}^3 \mapsto \mathbf{so}(3)$	<code>skew</code>
$\vee_\times(\cdot)$	unpack skew-symmetric matrix: $\mathbf{so}(2) \mapsto \mathbb{R}, \mathbf{so}(3) \mapsto \mathbb{R}^3$	<code>vex</code>
$[\cdot]$	augmented skew-symmetric matrix: $\mathbb{R}^3 \mapsto \mathbf{se}(2), \mathbb{R}^6 \mapsto \mathbf{se}(3)$	<code>skewa</code>
$\vee(\cdot)$	unpack augmented skew-symmetric matrix: $\mathbf{se}(2) \mapsto \mathbb{R}^3, \mathbf{se}(3) \mapsto \mathbb{R}^6$	<code>vexa</code>
$\text{Ad}(\cdot)$	adjoint representation: $\mathbf{SE}(3) \mapsto \mathbb{R}^{6 \times 6}$	<code>.Ad</code>
$\text{ad}(\cdot)$	logarithm of adjoint representation: $\mathbf{SE}(3) \mapsto \mathbb{R}^{6 \times 6}$	<code>.ad</code>
\circ	quaternion (Hamiltonian) multiplication: $\mathbb{H} \times \mathbb{H} \mapsto \mathbb{H}$	<code>*</code>
$\hat{\mathbf{v}}$	pure quaternion: $\mathbb{R}^3 \mapsto \mathbb{H}$	<code>Quaternion.pure</code>
\sim	equivalence of representations	
\simeq	homogeneous coordinate equivalence	
\ominus	smallest angular difference between two angles on a circle: $\mathbb{S}^1 \times \mathbb{S}^1 \mapsto \mathbb{R}$	<code>angdiff</code>
$\mathcal{K}(\cdot)$	forward kinematics: $\mathcal{C} \mapsto \mathcal{T}$	<code>fkine</code>
$\mathcal{K}^{-1}(\cdot)$	inverse kinematics: $\mathcal{T} \mapsto \mathcal{C}$	<code>ikine</code>
$\mathcal{D}^{-1}(\cdot)$	manipulator inverse dynamics function: $\mathcal{C}, \mathbb{R}^N, \mathbb{R}^N \mapsto \mathbb{R}^N$	<code>rne</code>
$\mathcal{P}(\cdot)$	camera projection function: $\mathbb{R}^3 \mapsto \mathbb{R}^2$	<code>.project</code>
\star	convolution	<code>iconv</code>
\otimes	correlation	
\equiv	colormetric equivalence	
\oplus	morphological dilation	
\ominus	morphological erosion	
\circ	morphological opening	
\bullet	morphological closing	
$\{F\}$	coordinate frame F	
$[a, b]$	interval a to b inclusive	
(a, b)	interval a to b exclusive, not including a or b	
$[a, b)$	interval a to b , not including b	
$(a, b]$	interval a to b , not including a	

MATLAB® Toolbox Conventions

- A Cartesian coordinate, a point, is expressed as a column vector.
- A set of points is expressed as a matrix with columns representing the coordinates of individual points.
- A rectangular region by two opposite corners $[x_{\min} \ x_{\max}; y_{\min} \ y_{\max}]$.
- A robot configuration, a set of joint angles, is expressed as a row vector.
- Time series data is expressed as a matrix with rows representing time steps.
- A MATLAB matrix has subscripts (i, j) which represent row and column respectively. Image coordinates are written (u, v) so an image represented by a matrix I is indexed as $I(v, u)$.
- Matrices with three or more dimensions are frequently used:
 - A color image has 3 dimensions: row, column, color plane.
 - A greyscale image sequence has 3 dimensions: row, column, index.
 - A color image sequence has 4 dimensions: row, column, color plane, index.

Common Abbreviations

2D	2-dimensional
3D	3-dimensional
DOF	Degrees of freedom
n -tuple	A group of n numbers, it can represent a point of a vector