

Introduction

Levente Tamás and Lucian Buşoniu

1 Brief Background

The field of robotics started already in the 19th century, with teleoperated vehicles. The motivation to further develop these devices arose from military interests, especially after the World Wars. The switch from remote controlled vehicles to autonomous ones began after the Second World War, when the early mobile robot called *Machina Speculatrix* was designed, which was able to follow a light source. The first boom in autonomous robotics was in the late sixties, and continues to the present day.

Robots are generally designed for transportation, manipulation, and surveillance tasks. Based on their configuration in space and the range of movement they can perform, one can distinguish between mobile robots (e.g. wheeled, underwater, or flying vehicles) (Ge and Lewis, 2006) and fixed object manipulators (Lewis et al, 2006). The mixture of these two types is usually referred to as mobile manipulators. For all these classes of robots, achieving autonomy crucially requires *automatic control*: algorithms that, without human assistance, are able to actuate the robot so as to achieve a desired configuration, to navigate through the environment, or to manipulate this environment in a useful way (Spong and Hutchinson, 2005; Lewis et al, 2006; Bruno and Oussama, 2008; Siegwart et al, 2011). Feedback from sensors is required since an exact model of the task is never available, and the robot must be able to compensate for model errors as well as unmodeled effects, such as a varying mass of the transported objects.

Traditionally, robot control deals with industrial robots, where the environment is predictable and the robot can function using models of the environment and pre-computed movements, with limited sensing. During the eighties the trend shifted from this classical way of thinking, dominant in the seventies, towards the reac-

The authors are with the Department of Automation, Technical University of Cluj-Napoca, Memorandumului 28, 400114 Cluj-Napoca, Romania, e-mail: levente.tamas@aut.utcluj.ro, lucian@busoniu.net

tive paradigm, which focuses more on sensor feedback (Brady et al, 1982). A further extension was the hybrid approach, using reactive principles at lower levels and higher-level model-based approximations (Khatib and Craig, 1989). More recently, the probabilistic robotics framework became dominant in research (Thrun et al, 2005). The probabilistic framework explicitly takes into account the inaccuracy in the models and sensors, and handles it in the control algorithms. This is important for robots to achieve autonomy outside the industrial setting, and to perform their tasks in uncertain, open environments. Sensing the environment is absolutely essential in this paradigm, and with new hardware such as stereo cameras, inertial units, and depth sensors, the autonomy of the robots is greatly expanded. High-speed and application-specific microprocessors enable the use of robots in real-time applications, by processing challenging large-volume sensing data from e.g. stereo cameras or depth sensors, and by allowing better control laws that take into account the complexity of the robot and environment dynamics.

The importance of robotic control is reflected by the focus placed on it in the top publication outlets in robotics on the one hand, and in systems and control on the other. For example, the International Conference on Robotics and Automation includes automatic control in the very title, and its 2014 edition included 6 workshops related to control; the same number was hosted by the 2014 International Conference on Intelligent Robots and Systems. The latest editions of the two main control events, the Conference on Decision and Control and the American Control Conference, dedicated specific tracks to control and sensing for (primarily mobile) robots. Robot control is also prominent in leading journals of the two fields: IEEE Transactions on Robotics, Robotics and Autonomous Systems, Automatica, Control Engineering Practice, etc.

Against this background, our book focuses on learning and sensing approaches to address the environment uncertainty, as well as on the control of networked and interconnected robots, as described next.

2 Goal and Motivation of the Book

While robots have long left factory floors, real penetration of advanced robotics outside the industry has been slow over the past decades, with research outcomes mainly remaining within the academia. The situation has however changed dramatically in recent years, with many novel marketable applications and robotic platforms appearing:

- domestic and assistive robots, such as Roomba, Mowbot, Create, and Aibo
- research and educational robots: Robotino, Mindstorm, PR2, TurtleBot
- surveillance in large, open environments for mapping, search & rescue, etc., with robots like the PackBot, Ranger, PatrolBot
- and of course the unprecedented explosion in unmanned aerial vehicles (UAVs) over the last couple of years, with proposed civilian applications ranging from

package delivery, through parking guidance, to delivering defibrillators to heart attack patients, see e.g. SUAS (2015)

Additional application domains are emerging, including surgical robots, surveillance in agriculture, space robotics, etc.

The defining characteristic of all these applications is the unpredictability and open, large-scale structure of the environment – which is often shared with humans. These features create several challenges for robot control, among which in this book we focus on two major ones. The first is *uncertainty* about the environment, coming either from the *limited sensors* available to measure the variables relevant for control or, more fundamentally, because the robot does not know how the environment *evolves and reacts* to its actions. Dealing with uncertainty is a traditional topic in robotics (Thrun et al, 2005; Stachniss, 2009) and overall in systems and control (Ristic et al, 2004), although it is still unsolved in general. In the absence of prior knowledge about the environment dynamics, learning a controller is the method of choice (Sutton and Barto, 1998; Sigaud and Peters, 2010; Lewis and Liu, 2012).

The second challenge we focus on is *networked structure*, which appears in many of the applications mentioned above. This is because the robot and its controller are often separated by a significant distance, while for mobile ground robots or UAVs wired connections are not feasible, and the robot must instead be controlled wirelessly. In all these cases, exchanging signals over a network is the best solution, but this comes with its own constraints and challenges that must be taken into account. Networked structure is also very important in large environments, where teams of robots are required (Balch and Parker, 2002) and communication among them to arrive at coherent sensor measurements or control actions is highly nontrivial. A particularly interesting problem occurs at the intersection of uncertainty and networked robotic teams: distributed sensing under uncertainty. When the uncertainty is driven by the physically distributed nature of the system, consensus algorithms can help in reaching an agreement on measurements. Overall, networked and interconnected systems comprise quite a new topic and their intersection with robotics is still in its starting phase, but the advance in the field has been quite significant in recent years (Shamma, 2007; Bullo et al, 2009; Bemporad et al, 2010).

For both classical and newer topics however, the recent growth in human-environment robotics has increased the pace of novel research. For the researcher or graduate student who is working in robot control, a resource is needed that presents recent advances on dealing with uncertainty and networked structure. Such a resource would also help researchers or PhD students who wish to enter robot control arriving from a related area (e.g., general control systems or computer vision). The aim of our book is to provide this resource: a snapshot of this area as it stands now, collecting in a single, coherent monograph a representative selection of state-of-the-art techniques. To this end, we have invited chapter contributions from experts in the relevant fields: robot control, learning control, state estimation, robot perception, and the control of networked and interconnected systems. To achieve a balanced viewpoint, we have included both already established, highly influential experts as well as younger researchers that have nevertheless already had a significant impact.

3 Book Structure

We structure the book along the three main challenges identified in Section 2 above: learning control to handle uncertainty about the dynamics; sensing under uncertainty, particularly as it pertains to control; and networked control of robots and multirobot systems. An overall view of this structure including chapter titles is given in Figure 1, and more details, including chapter outlines, are provided next.

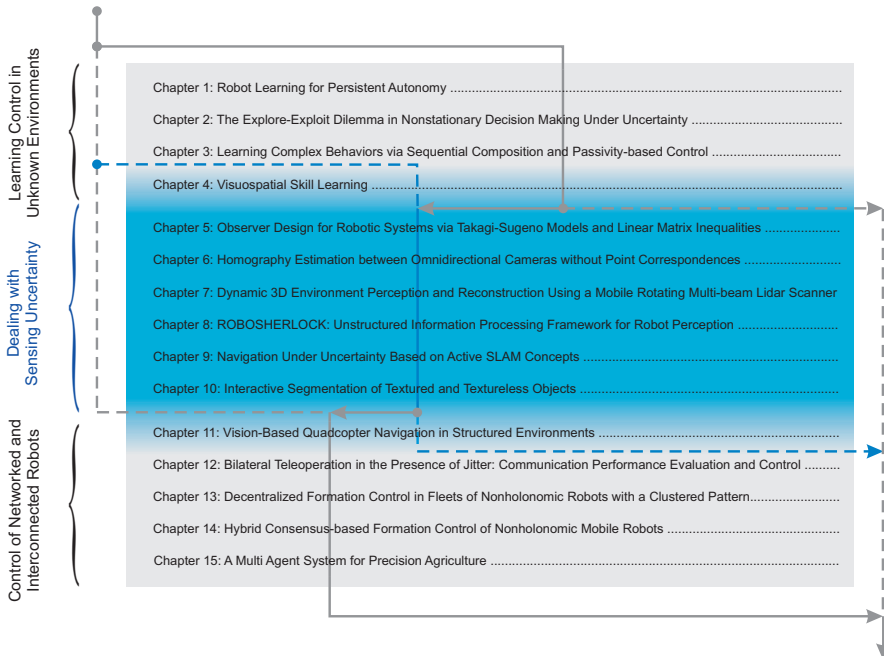


Fig. 1: Organization of the book. The background color changes for each main part, and the arrows indicate possible ways of reading the book.

The book starts with **Part I: Learning Control in Unknown Environments**, with a selection of learning-based techniques to handle uncertain dynamics. In order to develop robots that solve long-term missions in unknown, open environments, it is important to deal with failures (Chapter 1), to explore efficiently environments that change in time (Chapter 2), and to approach the problem in a practical way that exploits any available prior knowledge while learning to deal with the unknown parts of the environment (Chapters 3, 4). Starting off by imitating a human expert (Chapters 1, 4) is particularly promising.

In more detail, Chapter 1, *Robot Learning for Persistent Autonomy*, presents an overarching view and a significant step towards the major robotics goal of perform-

ing long-term autonomous missions. This is achieved by a combination of learning from expert demonstrations, and learning to recover from failures. Practical experiments illustrate the technique: valve-turning with the Kuka robot arm, and recovery from thruster failures with the Girona500 Autonomous Underwater Vehicle (AUV). Chapter 2, *The Explore-Exploit Dilemma in Nonstationary Decision Making Under Uncertainty*, again uses reinforcement learning methods in unknown environments but focuses on a complementary aspect: choosing when to exploit the information already gathered, versus exploration to learn more about the environment. This is done in the particularly challenging case of an environment that changes in time, and two methods are proposed to anticipate interesting changes. Simulated applications are given: planning least-visible paths for unmanned aerial vehicles (UAVs) in human environments, and surveillance via unattended ground sensors assisted by UAVs. Chapter 3, *Learning Complex Behaviors via Sequential Composition and Passivity-Based Control*, gives a modular approach where local controllers are learned with reinforcement learning and then sequentially composed with a finite-state machine, which is itself adaptive by finding the domains of attraction of each local controller. This idea is very useful for robots that operate in several modes, such as UAVs, which switch between takeoff, hovering, cruise flight, and landing modes. The approach allows including partial prior knowledge about the solution structure and dynamics, but does not require a full model – and learning tackles the unknown part. Chapter 4, *Visuospatial Skill Learning*, gives another modular approach that assumes predefined motor primitives such as grasping are available. Exploiting these primitives, learning is performed directly in the visual task space, starting from an expert demonstration and aiming to reproduce a given object configuration. The method is illustrated in simulations and on the Barret WAM robotic arm, which uses it to solve several real-life tasks.

The approach in Chapter 4 blurs the line between control and visual sensing, and so provides a transition to **Part II: Dealing with Sensing Uncertainty**. Even if the dynamics of the environment are fully known, before the robot can effectively solve a task there still remains the problem of finding the values of variables that are needed as inputs for control decisions. These variables are subject to uncertainty because the sensors of the robot cover a limited area, and extracting useful information often requires high-complexity processing of the raw data they provide (e.g. for stereo cameras or high density Lidars). We start by covering two basic problems for which research is still ongoing: determining the state variables (pose and velocities) of the robot itself (Chapter 5) and the relative poses of different cameras (Chapter 6). We then move on to higher-level perception methods for scene reconstruction and understanding (Chapters 7 and 8). We end Part II by two chapters that focus on the active sensing paradigm, which closes the loop between sensing and control in an interesting way, by *controlling* the robot so as to reduce the uncertainty in sensing. Active sensing is exploited to obtain better localization (Chapter 9) and to improve object segmentation (Chapter 10).

Specifically, Chapter 5, *Observer Design for Robotic Systems via Takagi-Sugeno Models and Linear Matrix Inequalities* deals with state estimation, which is made challenging by the nonlinearity of the dynamics. The mass matrix appearing in the

structure is exploited when representing the dynamics in a Takagi-Sugeno form, and the state estimator is proven to be convergent by Lyapunov techniques that boil down to solving linear matrix inequalities. A simulated example involving a two-wheeled mobile robot is provided. Chapter 6, *Homography Estimation between Omnidirectional Cameras without Point Correspondences*, presents a method to estimate the homography mapping between two omnidirectional cameras that look at the same scene. The method is novel in its use of matching segmented surfaces rather than pairs of points in the images, and its viability on real images is demonstrated.

Moving on to higher-level perception, Chapter 7, *Dynamic 3D Environment Perception and Reconstruction Using a Mobile Rotating Multi-beam Lidar Scanner*, describes a complete pipeline for online detection and tracking of pedestrians and vehicles, and for offline analysis of urban scenes, both from multi-beam LIDAR data. The method is evaluated in real urban environments. Chapter 8, *ROBOSH-ERLOCK: Unstructured Information Processing Framework for Robot Perception* describes an overall framework for perception that is able to respond to sensing queries from the controller phrased as high-level questions (such as ‘Where is object X?’). The framework was implemented in an open-source package and tested in a household environment using data acquired from a PR2 robot.

In Chapter 9, *Navigation Under Uncertainty Based on Active SLAM Concepts*, the objective is to control the robot so as to reduce the uncertainty in the robot’s location. The chapter provides an extensive overview of the active SLAM field, and describes a state-of-the-art approach that constructs a graph of robot configurations and then computes a minimal-uncertainty path using a graph search algorithm. The approach is validated on several public datasets. Chapter 10, *Interactive Segmentation of Textured and Textureless Objects*, goes beyond planning the trajectory, by allowing the robot to interact with the environment in order to reduce uncertainty. Specifically the robot grasps and moves objects in order to better segment them. Real-life evaluation is performed on cluttered scenes using the PR2 humanoid robot.

Part III: Control of Networked and Interconnected Robots deals with the network effects that appear for wirelessly controlled mobile robots and in communicating robot teams. This last part of the book starts with a transition Chapter, number 11, which still devotes significant attention to sensing but starts taking into account networked control effects. Chapter 12 studies the impact of such effects in robotic teleoperation. For the remainder of the book, we move to multirobot systems and tackle issues of agreement in the presence of communication constraints (Chapters 13, 14) and of cooperative control for a mixed ground-and-aerial robot team (Chapter 15).

Chapter 11, *Vision-Based Quadcopter Navigation in Structured Environments*, uses vision to detect and track the vanishing point of lines in perspective, in corridor and corridor-like environments. The robot (an AR.Drone 2 UAV) then navigates the environment by moving forward while keeping the vanishing point centered. Significant challenges arise due to the WiFi network interposed between the controller and the system: image frames arrive at varying intervals and many are dropped. This is compensated by running the filter in prediction mode. Chapter 12, *Bilateral Teleoperation in the Presence of Jitter: Communication Performance Evaluation and*

Control, studies in detail the impact of time-varying delays in a robotic teleoperation – still a single-robot system like in Chapter 11, but now operated through a haptic device. A controller designed in the passivity framework is proposed that can effectively deal with these delays. Practical experimental results are shown on a Kuka robot arm operated with the Sensable Phantom Omni haptic device.

In Chapter 13, *Decentralized Formation Control in Fleets of Nonholonomic Robots with a Clustered Pattern*, the objective is for a team of nonholonomic mobile robots to reach a given formation in a decentralized way. The robots are organized in subteams that communicate internally, and the team leaders communicate sporadically amongst themselves to achieve overall agreement. The network is limited to a graph structure due to limited range, both within the teams as well as between the leaders. Chapter 14, *Hybrid Consensus-based Formation Control of Nonholonomic Mobile Robots*, considers a similar problem of decentralized formation control, but with the addition that the robots must also navigate to a goal area. The robots form a single team (graph), and a control law is given that switches between two modes: formation-alignment, and navigating to the goal. Both Chapters 13 and 14 validate the techniques in simulations of mobile robot teams. Finally, Chapter 15, *A Multi Agent System for Precision Agriculture*, is focused on the emerging application of monitoring crops. A two-layer multirobot system is proposed, consisting of ground robots that navigate the cultivated field, and UAVs that act as longer-range sensors for the ground robots, by e.g. directing them to areas of interest and around obstacles. The components of the approach are tested in experiments on Surveyor SRV-1 ground robots and AR.Drone UAVs.

The book can be read in several ways, by following the arrows in Figure 1. Besides the default order of reading all the chapters in sequence (continuous line), the three parts are sufficiently self-contained to be read individually (dashed lines), with the following note. If the reader chooses to focus only on Part II, it is recommended to additionally read the connecting Chapters 4 and 11, since these contain important elements of sensing under uncertainty.

For the more application-oriented reader, an alternative way of looking at the book is by the type of robotic platform used in the experimental validation. Table 1 organizes the chapters by this criterion. Note that although they do not appear in this table, Chapters 3, 6, and 9 of course still provide experimental validations, but this validation does not concern a specific robotic platform since the method is general and can be applied to any platform (e.g. vision methods that are tested on public datasets).

Acknowledgment

The editors are grateful to the Romanian National Authority for Scientific Research, CNCS-UEFISCDI for supporting them financially through the Young Teams project No. PNII-RU-TE-2012-3-0040. We also wish to extend a deep thanks to the con-

Mobile ground robots	Chapters 5, 7, 13, 14, 15
Unmanned aerial vehicles	Chapters 2, 11 , 15
Autonomous underwater vehicles	Chapter 1
Robot arms	Chapters 1, 4, 12
Humanoid robots	Chapters 8, 10

Table 1: Chapters organized by the type of robotic platform used in the evaluations. Boldface font indicates real-life experiments, while the other chapters contain simulations.

tributing authors for their efforts, and to the editing team at Springer for their help throughout this project. Last but not least, we are very grateful to the anonymous reviewers of our book proposal for their comments, which led to significant improvements of the book.

References

- Balch T, Parker LE (eds) (2002) *Robot Teams: From Diversity to Polymorphism*. A K Peters / CRC Press
- Bemporad A, Heemels M, Johansson M (eds) (2010) *Networked Control Systems*. Lecture Notes in Control and Information Sciences, Springer
- Brady M, Hollerbach JM, Johnson T, Lozano-Perez T, Mason MT (1982) *Robot Motion: Planning and Control*. MIT Press
- Bruno S, Oussama K (eds) (2008) *Handbook of Robotics*. Springer
- Bullo F, Cortés J, Martinez S (2009) *Distributed Control of Robotic Networks. A Mathematical Approach to Motion Coordination Algorithms*. Princeton University Press
- Ge SS, Lewis F (eds) (2006) *Autonomous Mobile Robots: Sensing, Control, Decision Making and Applications*. CRC Press
- Khatib O, Craig JJ (1989) *The Robotics Review*. MIT Press
- Lewis F, Liu D (eds) (2012) *Reinforcement Learning and Adaptive Dynamic Programming for Feedback Control*. Wiley
- Lewis F, Dawson D, Abdallah C (2006) *Robot Manipulator Control: Theory and Practice*, 2nd edn. Control Engineering Series, Wiley
- Ristic B, Arulampalam S, Gordon N (2004) *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House
- Shamma J (ed) (2007) *Cooperative Control of Distributed Multi-Agent Systems*. Wiley
- Siegwart R, Nourbakhsh IR, Scaramuzza D (2011) *Introduction to Autonomous Mobile Robots*. MIT Press

- Sigaud O, Peters J (eds) (2010) From Motor Learning to Interaction Learning in Robots. Studies in Computational Intelligence, Springer
- Spong M, Hutchinson S (2005) Robot Modeling and Control. Wiley
- Stachniss C (2009) Robotic Mapping and Exploration. Springer
- SUAS (2015) URL <http://www.suasnews.com/>, Unmanned Aircraft System news website
- Sutton RS, Barto AG (1998) Reinforcement Learning: An Introduction. MIT Press
- Thrun S, Burgard W, Fox D (2005) Probabilistic Robotics. MIT Press