



Editorial

Robot learning from demonstration

1. Motivation

Programming by demonstration (PbD) is a key research topic in robotics. It impacts both fundamental research and application-oriented studies. Work in that area tackles the development of robust algorithms for motor control, motor learning, gesture recognition and visuo-motor integration. While the field existed for more than 20 years, recent developments, taking inspiration in biological mechanisms of imitation, have brought a new perspective. Programming by demonstration, now, encompasses more of the learning components of traditional approaches and is often referred to as *Learning from Demonstration* or *Learning by Imitation*. This issue aims at assessing the recent progresses in the field, gathering works following either an engineering or a biologically inspired approach.

2. Special issue composition

This special issue collects a subset of the best papers presented at the IROS'03 *Programming by Demonstration* workshop. The workshop took place on 31 October 2003 in Las Vegas, USA, at the occasion of the IEEE conference on Intelligent Robots and Systems (IROS'03). Eleven of the best papers presented at the workshop were subsequently submitted to peer review before publication in the special issue. The papers address a number of key questions underlying the PbD framework such as:

- What should the robot imitate, i.e. which features of the task should be reproduced?
- How could we define a general metric of imitation performance?

- How could the metric drive the choice of learning technique?
- What level of representation of movement is most suitable to both gesture recognition and motor control?
- Can a task/motion be decomposed into a set of primitives and how can one learn these primitives?
- How can models of human kinematics, used in gesture recognition, drive the reproduction of the task?
- When is prior knowledge required in order to ensure realistic timing for learning?
- How can knowledge be encoded to ensure its reuse in other learning tasks?

The papers have been ordered according to the general topic of PbD they tackle. These are summarized in each of the following subsections.

2.1. What to imitate

Billard et al. address the issue of discovering which features of the task should be reproduced and develop a formal mathematical framework to derive the metric of imitation performance and the optimal imitation control policy.

2.2. How to imitate

All other papers in this special issue address the core problem of how to imitate, that is, how to recognize and encode human motion in a way that makes it easily transferable to a robot. These papers have been ordered according to the level of imitation they tackle, from imitation of low-level features, namely

joint trajectories, to imitation of higher level features, such as imitation of complete actions, tasks and behaviors.

2.3. Imitation of trajectories

Nakanishi et al. and Ude et al. develop methods for deriving robust imitation of joint trajectories. Nakanishi et al. tackle imitation learning of locomotion patterns for a bipedal robot using Locally Weighted Regression as learning scheme, a gradient-descent based method. Ude et al. tackle imitation of full body motion by a 30 DOFs humanoid robot and develop an optimized method for deriving the segmentation of trajectories into B-spline wavelet functions.

2.3.1. Imitation of actions/tasks

Dillmann; Zhang and Rossler; Chella et al.; Steil et al.; Aleotti et al.; Bentivegna et al. develop methods to recognize and to imitate complex manipulatory tasks. These works stress the need of introducing prior knowledge in the way information is encoded to achieve fast and reusable learning.

Dillmann builds upon the long-standing PbD work done in his laboratory and lists a number of key problems to be addressed in order to build a system capable of learning a wide range of housing tasks. Steil et al. and Zhang and Rossler develop comprehensive physical systems for recognizing and reproducing on-line sequences of manipulative actions. Chella et al. focus on robust recognition of hand postures for imitation of meaningful gestures, while Aleotti et al. focus on recognizing finger motion for imitation of fine manipulation in pick-and-place tasks. Bentivegna et al. develop a metric-based method for determining the most likely strategy to follow while playing air hockey or while wandering in a maze.

2.3.2. Imitation of behaviors

Le Hy et al.; Wermter et al. address the issue of recognizing and imitating behaviors. Behaviors consist, here, of appropriate sensory-mapping. Wermter et al. position paper offers a biologically-inspired connectionist approach to learning the appropriate auditory-visuo-motor mapping. Le Hy et al. use a Bayesian framework to determine the most likely sensory-motor transition, based on a list of potential sensory-motor transition.

3. Conclusion

This issue is not a comprehensive overview of the large field of Robot Programming by Demonstration, but it conveys a summarized view of the state-of-the-art approaches to the field. While the field has been successful at solving a large number of issues of PbD, a number of fundamental questions have yet to be tackled. These are: *Does imitation speed up skill learning in robots?* If yes, *What are the costs of imitation learning?* More fundamentally: *Are there skills and tasks that cannot be acquired without demonstration?* And, conversely: *Are there skills and tasks that cannot be learned by imitation?* Finally, *Can imitation use known motor learning techniques or does it require the development of new learning and control policies?*

As the field is constantly and quickly growing, we give you rendez-vous in a few years to review the field's progresses in a similar special issue.



Aude Billard is an SNF Professor at the School of Engineering at the Swiss Federal Institute of Technology (EPFL). She is also an Adjunct Assistant Professor at the Computer Science Department at the University of Southern California. She received her BSc (1994) and MSc (1995) in Physics from the EPFL, with specialization in Particle Physics at the European Center for Nuclear Research (CERN). She received her MSc in Knowledge Based System (1996) and her PhD in Artificial Intelligence (1998) from the Department of Artificial Intelligence at the University of Edinburgh. She, then, was a Post-doctoral Fellow at IDSIA and LAMI (EPFL, 1998–1999). She was a Research Associate (1999–2000) and a Research Assistant Professor (2000–2002) at the Computer Science Department at the University of Southern California. Her research interests cover the fields of Artificial Neural Networks, Robotics, Neural Modeling, Computational Neuroscience and more generally Machine Learning. Her work tackles special topics, such as Programming Through Demonstration, Imitation Learning and Language Acquisition.



Roland Siegwart (1959) received his MSc ME in 1983 and his Doctoral degree in 1989 at the Swiss Federal Institute of Technology (ETH) Zurich. After his PhD studies he spent 1 year as a postdoc at Stanford University where he was involved in micro-robots and tactile gripping. From 1991 to 1996 he worked part time as R&D director at MECOS Traxler AG and as lecturer and deputy head at the Institute of

Robotics, ETH. Since 1996 he is a full professor for Autonomous Systems and Robots at the Swiss Federal Institute of Technology, Lausanne (EPFL), and since 2002 also vice-dean of the School of Engineering. He leads a research group of around 25 people working in the field of robotics and mechatronics. He published over 100 papers in the field of mechatronics and robotics, is an active member of various scientific committees and co-founder of several spin-off companies. He was the general Chair of IROS 2002 and he is currently VP for Technical Activities of the IEEE Robotics and Automation Society (<http://asl.epfl.ch/>).

Aude Billard*

Roland Siegwart

ASL3-I2S-STI, EPFL, 1015 Lausanne, Switzerland

*Corresponding author

E-mail addresses: aude.billard@epfl.ch (A. Billard)

roland.siegwart@epfl.ch (R. Siegwart)