Pattern Recognition Letters 67 (2015) 91–99

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Cognitive system for autonomous underwater intervention☆

Arnau Carrera[a,*], Narcís Palomeras[a], Natàlia Hurtós[a], Petar Kormushev[b], Marc Carreras[a]

[a] Underwater Vision and Robotics Research Centre (CIRS), University of Girona, c/Pic de Peguera, 13 (Parc Científic), 17004, Girona Spain
[b] Istituto Italiano di Tecnologia, via Morego, 30, 16163, Genova Italy

ARTICLE INFO

Article history:
Available online 18 June 2015

Keywords:
Learning by demonstration
Imitation learning
Underwater intervention
Autonomous underwater vehicles
Cognitive system
Knowledge acquisition

ABSTRACT

The implementation of autonomous intervention tasks with underwater vehicles is a non-trivial issue due to the challenging and dynamic conditions of the underwater medium (e.g., water current perturbations, water visibility). Likewise, it requires a significant programming effort each time that the vehicle must perform a different manipulation operation. In this paper we propose, instead, to use a cognitive system that learns the intervention task from an expert operator through an intuitive learning by demonstration (LbD) algorithm. Taking as an input few operator demonstrations, the algorithm generalizes the task knowledge into a model and is able to control the vehicle and the manipulator simultaneously to reproduce the task, thus conferring a more adaptive behavior in front of the environment changes and allowing to easily transfer the knowledge of new tasks. A cognitive architecture has been implemented in order to integrate the LbD algorithm with the onboard sensors and actuators and to allow its interplay with the vehicle perception, control and navigation modules. To validate the full framework we present real experiments in a water tank using an AUV equipped with a four DoF manipulator. A human operator teaches the system to perform a valve turning intervention and we analyze the results of multiple task reproductions, including cases under the effect of water current perturbations, showing the success of the system in autonomously reproducing the task.

1. Introduction

Cognitive systems hold the promise for changing future human–machine interactions through integrated systems that are able to perceive, understand, learn and develop in their environment [26]. Knowledge representation, and particularly how this knowledge is transferred between humans and machines becomes therefore a topic of utmost importance for cognitive systems. A common option to represent and transfer knowledge is the combination of natural language processing with machine learning techniques [9]. However, if we wish to transfer the ability to perform a task or a skill to a machine, natural language might become too intricate, while providing the machine with an example of the task itself can be a much more straightforward way to convey the associated knowledge. For these cases, learning by demonstration (LbD) algorithms [2] offer a natural solution for transferring the knowledge of a new skill to a system by extracting it from a set of expert demonstrations. Opposite to other demonstration-based systems, where the robot learns to accomplish a task by naively recording predefined paths or trajectories and reproducing them later, LbD extracts the knowledge embedded in multiple demonstration trajectories and generalizes them into a representative model of the task, that is later used to generate reproductions. The main advantage of this paradigm is that it confers a more adaptive behavior in front of environment changes in addition to increase flexibility and scalability as the system can learn new tasks without requiring additional programming. Hence, learning by demonstration can be an intuitive and effective way to build a cognitive model and develop a cognitive system that can perceive, learn and actuate on its environment in an autonomous way.

While these kind of algorithms have been applied in teaching some tasks to robotic manipulators [3], our aim in this paper is to apply an LbD technique to an autonomous underwater vehicle (AUV) to enable it to learn, intuitively, a sub-sea intervention task. Underwater manipulation operations are often required, for instance in the maintenance of permanent sub-sea observatories, deployment and recovery of benthic stations, or in the inspection, repair and maintenance of submerged infrastructures of the offshore industry. These sort of intervention tasks are nowadays carried out using remotely operated vehicles (ROVs), that require expensive support vessels and dedicated crew and operators. In this sense, performing them autonomously with an AUV endowed with a manipulator (the so-called Intervention-AUVs (I-AUVs)) would convey significant advantages in terms of operational time and costs. The implementation of autonomous intervention tasks with these vehicles is a non-trivial issue due to the challenging and dynamic conditions of the underwater

☆ This paper has been recommended for acceptance by Lledó Museros.

* Corresponding author. Tel.: +34 972 419 651, +34647154368.
E-mail address: arnau.carrera@udg.edu, arnaucarrera@gmail.com (A. Carrera).

http://dx.doi.org/10.1016/j.patrec.2015.06.010
0167-8655/

medium: water perturbations (i.e., current, waves), reduced visibility, difficulties in understanding the scene, and a high degree of uncertainty in navigation and perception sensors. Likewise, it requires a significant programming effort each time that the vehicle must perform a different manipulation operation. For those reasons, we believe that approaching the problem using a LbD algorithm can contribute to advance the state of the art of autonomous underwater intervention.

It is worth noting that in the last years some research projects have begun to demonstrate autonomous intervention capabilities although none of them has used machine learning techniques but classical manipulation theory. The SAUVIM project [16] proposed an underwater intervention using a priority order controller for the manipulator to recover objects of the seafloor. In the TRIDENT project [21] a system to search and recover objects with a light I-AUV was presented. The I-AUV implements an underwater vehicle manipulator schema (UVMS) which is guided by a visual system to grasp the object. The TRITON project [8,19] shows some manipulation with an I-AUV docked in a sub-sea panel. The manipulation is guided to follow a trajectory composed by several points generated from the visual detection of the valve. The work presented here is in the context of the PANDORA [10], whose aim is to increase the persistent autonomy in underwater operations. It constitutes, to the best of our knowledge, the first application of a machine learning technique on an I-AUV to learn an intervention task, in this case performing a valve turning manipulation in free-floating mode.

The application of LbD techniques in the underwater domain presents several added complications due to the previously mentioned characteristics of the underwater medium. For this reason, the implementation of the LbD algorithm has required to develop a complete framework that combines navigation, control, and perception to fully integrate the learning algorithm within the vehicle cognitive architecture. The proposed framework, built on our previous work [6], is designed to learn eight degrees of freedom (DoF) to control simultaneously the trajectory of an AUV and its manipulator. To perceive the environment, the system uses the vehicle cameras and also a force and torque (F/T) sensor to detect the contact between the manipulator and the target object. Information from all these sensors is acquired while a pilot is performing a demonstration of the intervention task to be learned. Then, from a set of several demonstrations, the proposed LbD algorithm generalizes a control policy able to accomplish the intervention task with the same performance than the human operator. Moreover, the previously presented framework has been complemented with a fuzzy decision maker (RFDM) algorithm that evaluates the risk of performing the manipulation according to the environment conditions [1] and a partial-order task planner [7] that manages the high-level decision making of the intervention mission.

To validate the proposed approach the implemented framework has been tested in the context of a valve turning intervention experiment where an I-AUV was set to turn the *t*-shaped valve handles of a sub-sea panel mock-up. We have further extended the results presented in [6] to demonstrate the persistent operation of the developed cognitive system in a long experiment conducted in a water tank during more than 3 h. In this experiment the vehicle has been periodically localizing the panel and turning its valves to different configurations along the time according to a predefined plan. Besides, to simulate more realistic conditions and test the adaptability of the manipulation trajectory under perturbations we have introduced water currents through two external thrusters. Results prove the good performance of the involved learning, control, perception and planning techniques even under the presence of high currents.

The rest of this paper is organized as follows. Section 2 overviews related work on LbD in the robotics community and describes the LbD algorithm that has been implemented. Section 3 describes the vehicle used to perform the intervention task as well as the developed intervention framework. Results obtained from the long valve turning mission are presented and analyzed in Section 4. Section 5 summarizes, and concludes the work.

2. Knowledge acquisition by means of learning by demonstration

In the context of cognitive systems, natural language is among the most popular choices to represent knowledge. Using this representation, it is possible to reason about large amounts of data and extract relevant information [24,25]. However, for describing the actions to be performed by an I-AUV, the kinesthetic representation of these actions can be more simple and efficient than semantically describing the task at hand. Along this idea, the LbD machine learning technique has opened new avenues to transfer knowledge from an expert human operator to a machine through demonstrations.

This type of algorithm follows three sequential phases: first, a set of *demonstrations* of the task are recorded; second, the algorithm *learns* by generalizing all demonstrations and creating a model; finally, the algorithm loads the model and uses it to *reproduce* the task.

2.1. LbD related work

Different LbD algorithms have been proposed throughout the literature, depending on the method used to encode the learned trajectory. Calinon et al. [5] proposed a representation based on Gaussian mixture model (GMM), which was later extended by Krüger et al. [15] using incremental GMM to automatically set the number of Gaussians. Furthermore, Calinon et al. [4] parameterize the model with the relevant coordinate system of the task and adapt the model in real-time to changing position of the relevant elements in the environment. Similar to the GMM, a hidden Markov model (HMM) [12] has also been used to represent a trajectory, together with its correspondent parametrized version by Kruger et al. [14]. Both GMM and HMM representations require the use of a regression algorithm like the Gaussian mixture regression (GMR) to generate a desired trajectory with an associated density distribution.

A different approach is to use dynamic movement primitives (DMP) [11,20]. Unlike GMM and HMM, DMP uses the learned model to dynamically generate the required commands to perform the reproduction of the trajectory. This makes the approach more robust to external perturbations and easily adaptable to different domains. DMP has also been parameterized by Matsubara et al. [17] and extended by Kormushev et al. [13] to include a force associated with the trajectory.

Therefore, given the simplicity of the representation and its flexibility, DMP is more suitable in the context of this work and has been chosen as the base of our learning framework.

2.2. Dynamic movement primitives (DMP)

DMP is an algorithm where the learned skill is encapsulated in a superposition of basis motion fields (see Fig. 1). The method used in this paper is an extension of the DMP proposed by Kormushev et al. [13]. The flexibility of the representation allows the adaptation of the algorithm to specific requirements, as it will be described in Section 2.3.

To better understand this encoding, we can imagine a mass attached to different damped strings. These strings attract the mass changing their forces along the time of the experiment, moving the mass following the desired trajectory.

To generate the superposition each attractor has an associated weight which changes along the time defined by the $h_i(t)$ function (1). The weight of each attractor is represented with a Gaussian, whose centers μ_i^T are equally distributed in time, and whose variance parameters $\Sigma_i^T = \text{total\_time}/K$ are set to a constant value