



A cognitive architecture for automatic gardening



Alejandro Agostini^{a,*}, Guillem Alenyà^b, Andreas Fischbach^c, Hanno Scharr^c, Florentin Wörgötter^a, Carme Torras^b

^a Bernstein Center for Computational Neuroscience, Göttingen, Germany

^b Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Barcelona, Spain

^c Institute of Bio- and Geosciences: Plant Sciences (IBG-2), Forschungszentrum Jülich GmbH, Jülich, Germany

ARTICLE INFO

Article history:

Received 7 July 2016

Received in revised form 21 April 2017

Accepted 22 April 2017

Keywords:

Cognitive architecture
Optimized plant treatments
Automatic gardening
Human-robot interaction
Learning planning operators

ABSTRACT

In large industrial greenhouses, plants are usually treated following well established protocols for watering, nutrients, and shading/light. While this is practical for the automation of the process, it does not tap the full potential for optimal plant treatment. To more efficiently grow plants, specific treatments according to the plant individual needs should be applied. Experienced human gardeners are very good at treating plants individually. Unfortunately, hiring a crew of gardeners to carry out this task in large greenhouses is not cost effective. In this work we present a cognitive system that integrates artificial intelligence (AI) techniques for decision-making with robotics techniques for sensing and acting to autonomously treat plants using a real-robot platform. Artificial intelligence techniques are used to decide the amount of water and nutrients each plant needs according to the history of the plant. Robotic techniques for sensing measure plant attributes (e.g. leaves) from visual information using 3D model representations. These attributes are used by the AI system to make decisions about the treatment to apply. Acting techniques execute robot movements to supply the plants with the specified amount of water and nutrients.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

While agriculture tasks entailing repetitive action on the basis of immediate sensing have been successfully automated (Blackmore, 2007), gardening considering longer time frames has remained elusive. Taking individualized care of plants is difficult because although plants react to changes of light, water and/or nutrients, these reactions come with long, variable delays and are also quite variable in strength. Moreover, individual reactions are history dependent, because an initially healthy plant will react differently to a treatment than a stressed or damaged one.

Human gardeners take actions relying on their *expertise* and their knowledge of the *history of events* observed in a given plantation. However, to attend to the needs of a plantation in large industrial greenhouses, having a team of gardeners might not be cost effective. The alternative is to use a robotic system capable of taking actions according to the specific needs of every plant. This is the main focus of this paper, where a robotic cognitive architecture for automatic gardening is presented. The architecture integrates artificial intelligence techniques for decision-making with robotics

techniques for sensing and acting. Low-level robotic approaches comprises short-term interaction of the robot with plants entailing 3D model acquisition of deformable objects (leaves) (Alenyà et al., 2013) and robot-arm manipulation approaches for active vision and individual treatment application (Fischbach, 2011).

For making decisions we implemented the decision-making framework presented in Agostini et al. (2014), specifically designed to generate action plans in tasks involving long delays between actions and effects. This framework focuses on the cognitive abilities needed, namely decision-making based on logic-based planning and learning weakly correlated cause-effects along sequences of events. These mechanisms compile expertise so as to deduce, from past sequences of events, the long-term treatments producing the best desired results in each particular situation of a plant. Planning and learning are interleaved in a way that permits uninterrupted operation by requesting a human gardener to instruct the treatments when the knowledge acquired so far is insufficient to make a decision. In this case, the gardener simply specifies the immediate treatment to apply to the plant while the learning approach autonomously finds the relevant events that compactly describe the plant evolution under that treatment. For instance, the gardener may instruct the robot to provide small doses of nutrient in combination with doses of water

* Corresponding author.

E-mail address: aagosti@gwdg.de (A. Agostini).

three times a day, if he detects a low growth rate of the plant due to lack of nutrients, or simply small doses of water twice a day, if he knows that large doses of nutrients and water has been administered right before. With these instructions, the system progressively improves the decision-making performance by coding plant evolutions into planning operators until the human intervention is not longer required.

Many industrial and service tasks require these cognitive abilities, and we have worked on the gardening setting because of our involvement in the European project GARNICS (GARNICS, 2010–2013). The goal of GARNICS was to automatically monitor large botanic experiments in a greenhouse to determine the best treatments (watering, nutrients, light) to optimize predefined aspects (growth, plant status) and to eventually guide robots to obtain the required measures from leaves and apply the treatments prescribed by the system.

The cognitive architecture has been tested in an industrially-relevant plant-phenotyping application (Houle et al., 2010; Furbank and Tester, 2011), with very encouraging results as presented in Section 3. In the final experiment within the GARNICS project, images of Tobacco plants (*Nicotiana tobacum* cv. Samsun) were acquired every hour for up to 30 days, from which the required appearance features were extracted. Moreover, for each individual plant, light intensity was measured, and water and nutrients were dispensed using an automated flexible-tube pump. The image dataset has been made publicly available (Scharr et al., 2014; Minervini et al., 2015) and further details on how these experimental data were collected can be found therein (cmp. also Section 3, esp. Fig. 6).

1.1. Related works

Robotics applications for the execution of human-like tasks have been tackled using techniques of human-robot interaction, task planning, and symbol grounding (Ingrand and Ghallab, 2014). Artificial intelligence planning techniques play a fundamental role in such applications since they use a symbolic, logic-based notation compatible with human language. This allows for a natural human-robot interaction, letting a lay person easily provide instructions to the robot (Argall et al., 2009). For a successful integration of these techniques in a robotic platform, it is mandatory to ground the symbolic descriptions of the AI planning methods to let the robot interact with the real world in order to execute a task (Harnad, 1990). This is normally done by integrating methods of different levels of abstractions (Ingrand and Ghallab, 2014; Krüger et al., 2011; Beetz et al., 2010; Dantam et al., 2016; Paxton et al., 2016), ranging from purely symbolic methods (Ghallab et al., 2004), to sensor information processing (Szeliski, 2010) and acting methods (Kemp et al., 2007). Several approaches integrate planning and acting for the robotic execution of human like tasks. In Dantam et al. (2016), a new architecture that integrates task and motion planning (TMP) is proposed. This architecture uses an incremental constraint-based task planning to dynamically incorporate motion feasibility at the task level during task execution, facilitating the symbol grounding problem. The architecture presented in Dianov et al. (2016), on the other hand, incorporates human assistance to demonstrate relevant actions for the task. The main novelty of their approach is a method for learning from demonstration that permits reusing knowledge of previously learned tasks to accelerate the learning of new ones. This is done by exploiting semantic similarities between tasks parameters. Learning from demonstration has also been used in the architecture presented in Paxton et al. (2016). In this case, a human provides example executions of symbolic actions that are used to update low-level probabilistic models parametrized for

each specific robotic platform, which permits evaluating the feasibility of grounding symbolic actions.

Robotics applications has also been implemented for the task of automatic gardening (Correll et al., 2010; Blackmore, 2007; Al-Beeshi et al., 2015). In Correll et al. (2010) a set of mobile robots are equipped with eye-in-hand cameras and mobile arms to water and harvest tomato plants according to their individual needs. To recognize position and maturity level of tomatoes (green or red), the system selects a robot that moves to the plant to take images from six (fixed) different perspectives. Tomatoes are recognized using a feature-based method that detects circles and smooth areas using a convolution approach (Torralba et al., 2004). A humidity sensor is placed in the soil to evaluate watering requirements. Task planning is carried out to allocate plants to robots. However, they do not apply task planning or any other AI approach to decide the treatment for each plant according to its needs. Water is supplied to the plant in a fixed amount and in a reactive manner when the value of the humidity sensor drops below a predefined threshold. Another example of a robotic system applied to automatic gardening is presented in Al-Beeshi et al. (2015). In this case, the system is split in two parts: one part attached to the greenhouse for the control of the environmental light, temperature, and humidity, and the other one fixed on a robot to control plant watering and seeding. Watering and seeding take place at predefined positions in the soil and are triggered by manually activated buttons. As in Correll et al. (2010), humidity sensors are placed in the soil to reactively water plants using a thresholded approach. In their architecture, none of the automatic processes uses artificial intelligence techniques for planning or learning.

In the following sections we present our cognitive architecture for automatic gardening. Section 2 introduces the general architecture of the system. In this section the robotic mechanisms for sensing and acting as well as the artificial intelligence techniques for planning and learning are described. The performance of the cognitive architecture in a real scenario is assessed in Section 3 and deeply discussed in Section 4. The paper ends with some conclusions (Section 5).

2. General architecture of the system

Fig. 1 presents a general scheme of the cognitive architecture. The architecture is composed of a decision-making framework (DMF), in charge of deciding the action (treatment) to be applied to the plant, the modules for perception and action, which ground the symbolic states and actions, respectively, and a KUKA arm robotic platform equipped with a set of cameras and a probe to supply water and nutrients.

The DMF receives a description of the plant from the *perception module*, which transforms the images captured by the cameras into a symbolic description of the plant state (e.g. number of leaves, size of the plant, etc.). The state description is stored in the *history* database that is used by the *decision maker* to define which treatment to apply depending on the previous history of the plant and the goal specification. The treatment is then sent to the *action module* for its actual execution. The history database is also used to learn which of the past experienced events are actually relevant to predict the evolution of the plant under a treatment. This is the task of the *learner*, which generates and refines *planning operators* (POs) containing these relevant events. The system is supported by a human *gardener* that specifies the treatment to apply when the POs learned so far are not enough to make a decision in the current situation. After the treatment execution, the process starts over again by generating a new plan from the reached plant state. This is a strategy known as replanning, which enables the robot to immediately make use of the new knowledge acquired after the

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات