A Cognitive Approach to Affordance Learning in Robotic Ecologies

Mauro Dragone(✉)

Trinity College Dublin, College Green, Dublin 2, Ireland
dragonem@tcd.ie
https://www.scss.tcd.ie/~dragonem/

Abstract. The Robotic Ecology vision shares many similarities with the one pursued by the IoT community: The ideal aim on both fronts is that arbitrary combinations of devices should be able to be deployed in unstructured environments, such as those exemplified in a typical household, and there efficiently cooperate to the achievement of complex tasks. While this has the potential to deliver a range of modular and disruptive applications, a pressing and open research question is how to reduce the amount of pre-programming required for their deployment in real world applications. In order to inspire similar advancements within the IoT community, this extended abstract discusses how this goal has been addressed by pioneering the concept of a self-learning robotic ecology within the EU project RUBICON (Robotic UBIquitous Cognitive Network); how such an approach relates to the concept of Affordances at the basis of Gibsons' theory of ecological psychology; and how it can be used to drive the gradual adaptation of a robotic ecology to changing contexts and evolving requirements.

Keywords: Robotic ecology · Affordances · Cognitive systems

1 Robotic Ecologies

In [1], Saffiotti and Broxvall discuss the implications of their PEIS-Ecology instantiation of the Robotic Ecology approach from an ecological (Gibsonian) point of view, by conceiving the interaction between each device and its environment in terms of mutuality and reciprocity. An ecology of simple devices can achieve complex tasks by performing several steps in a coordinated fashion while also exchanging sensor data and other useful information in the process. Note how such a viewpoint falls into the framing assumption of ecological psychology, which, as Greeno notes [2] "*... involves a shift of the level of primary focus of cognitive analysis from processes that can be attributed to individual agents to interactive processes in which agents participate, cooperatively, with other agents and with their physical systems that they interact with*".

Saffiotti and Broxvall emphasise how its embodied nature is what makes confronting an ecological view in a robotic ecology characteristically different from what is usually done in pure software systems, e.g., for the orchestration

R. Giaffreda et al. (Eds.): IoT360 2014, Part I, LNICST 150, pp. 54–57, 2015.
DOI: 10.1007/978-3-319-19656-5_8

of web-services. Noticeably, the semantic of services (the function that they can provide) is usually advertised to the network irrespective of the components that may actually come to use them. On the contrary, just as with Affordances [3], intended as behavioural possibilities, a robotic ecology must be described in terms of relations (configurations) between its components. For instance, a ceiling camera can track and inform a mobile robot about its location in the environment (thus *affording localization* to the robot). However, this can be done as long as the robot in question (i) is not too small, (ii) is inside the field of view of the camera, and (iii) its color can be distinguished from the one of the floor.

A PEIS Ecology adopts temporal constraints to represent the possible relations between PEIS components, e.g., stating that navigation must occur while location information is available, or that a class of robots is able to push a given object. These relations are then shared and resolved into global configurations, by using either a centralized constraint satisfaction solver or a multi-agent, reactive approach. However, both methods lack learning capabilities and rely instead on pre-programmed, static and brittle domain knowledge. The inability of modelling other (unforeseen at design time) information in the domain, including possible synergies, conflicts and other inter-dependencies between their many components, is what breaks the modularity of the development of these systems, and ultimately impacts on their ability to pro-actively and smoothly adapt to changing contexts and evolving requirements. These are the key issues addressed in RUBICON [4], as outlined in the next section.

2 The RUBICON Approach

The most appealing aspects of the concept of affordances as a source of inspiration in Robotics are (i) its implicit emphasis on the relationship between an agent and the environment, and (ii) its grounding in the paradigm of direct perception. The central question for Gibson was whether affordances can be directly perceived. In their stride to build physically embodied agents, roboticists have thus sought to enable robots with the ability to learn to recognize affordances, thus ultimately reducing the complexity of representation and reasoning.

A common approach has been the utilization of an exploratory stage, in which the robot tries out different action primitives and observes their consequences. To this end, roboticists have usually based their works on more formal, e.g., functional elaboration of affordances seen as opportunities for action and inherently suited to be used as pre-conditions in a planning context by virtue of their predictive quality (see [5] for a survey). A service robot may thus learn how to poke, push, pull, rotate and lift objects, and also for what objects and in what situations its actions can more successfully achieve given results. A mobile robot may learn to infer the traversability of its surroundings by mapping from the space of the features extracted from its range sensors, to the effects (success/fail) of a number of basic manoeuvres. Techniques supporting online learning usually exploit curiosity measures to guide the robot's exploration process, by reducing unnecessary interaction with the environment when the robot is confident that it will not bring about new information.

Equipping robotic ecologies with similar learning abilities poses a formidable number of issues: from the computational constraints and the number of the devices involved, to the difficulty of identifying suitable and reliable teaching information to drive system's adaptation. RUBICON [4] addressed these problems by supporting a self-sustaining dynamic between cognitive capabilities realized in a modular architecture, shown in Fig. 1.

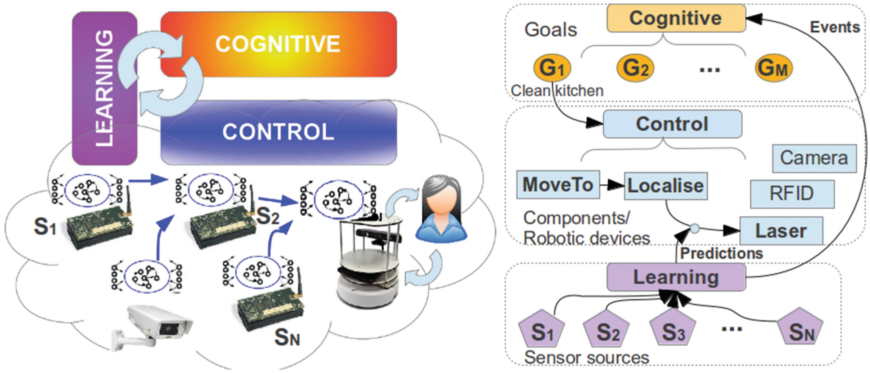

Fig. 1. High-level, hierarchical RUBICON architecture: Sensor data is processed as much as possible locally on computational constrained and robotic devices. Information is extracted and exploited by the higher layers.

The central component of the RUBICON architecture is a plan-based Control Layer (an evolution of the system employed in PEIS, as described in Sect. 1), which is able to decide which components (e.g. robot behaviours) and/or devices need to collaborate to achieve given service goals. The key approach to enabling system adaptation is to (i) improve its ability to extract meaning from noisy and imprecise sensed data, and (ii) learn what goals to pursue, and (iii) how to pursue them, from experience, rather than by relying on pre-programmed goal-deliberation strategies and plan pre-conditions.

The first and the last of these challenges are met by the Learning Layer, a distributed and adaptable learning infrastructure based on echo state networks, and which is used to process time series of data gathered by the sensors in the ecology. Its outputs are used to recognize events concerning the state of the environment and of the users, and predict the success/fail rate in using given devices and/or other components in the ecology. This information is used by the Control Layer to inform its configuration of the ecology. The second challenge is the responsibility of the Cognitive Layer, which uses Self-Organising Fuzzy Neural Networks to reason over the events recognized by the Learning Layer in order to predict the user's activation of appliances and robotic services.

In our reports [4], we have described how our system, equipped with some initial knowledge and supervised information, can be trained to provide some basic services. The system can use this as a starting point and self-adapt to the preferences of the user and to modification to the environment in a number of ways: firstly, the Cognitive Layer can learn to predict the user's routines. For instance, by observing past instances in which the user has summoned her cleaner robot to the kitchen after eating her meal, the system learns to send the robot without waiting for the user's request. Secondly, the Control Layer can monitor the performances and the outcomes of its own plans, and feed them (as teaching signals) to the Learning Layer in order to learn previously unmodelled plan pre-conditions. For instance, in [6], we have shown how a robot can learn to use the information it receives from itw own sensors and the sensors embedded in the environment (e.g. infrared sensors signalling the movements of the user and/or the robot) to predict (i) what are the best situations in which cleaning a certain room will be less likely to annoy the user, and (ii) when is best to use the RFID-based localization component in place of its own laser (e.g., after the user installs a new mirror that disturbs the robot's laser). Finally, the Cognitive Layer can "explore" (in a manner similar to the one adopted in curiosity-driven, robot manipulators) by trying out new goals in different situations in order to gather further experience and/or feedback from the user.

Our approach allows robotic ecologies to be driven by easily identifiable (albeit rough) rules, while delegating, over time, symbolic reasoning to data-driven inference for the purpose of learning to recognize the affordances of their environment, directly from sensor features. This is a clear improvement over past solutions, which demanded for all goal rules and plan pre-conditions to be specified a priori. What makes it a practical approach, which limits the computational cost of our learning solutions and enables us to use fully automatic feature selection algorithms, is (i) the existence of a finite set of pre-defined and simple goals, (ii) their clear distinction to the plans to achieve them, and (iii) the reduced number of sensor sources included in current smart homes. Future research should increase the scalability of these systems, and address the challenging problem of autonomously learning what goals are achievable to the ecology.

References

1. Saffiotti, A., Broxvall, M.: Affordances in an ecology of physically embedded intelligent systems. In: Rome, E., Hertzberg, J., Dorffner, G. (eds.) Towards Affordance-Based Robot Control. LNCS (LNAI), vol. 4760, pp. 106–121. Springer, Heidelberg (2008)
2. Greeno, J.G.: Gibson's affordances. Psychol. Rev. **1012**, 336–342 (1994)
3. Gibson, J.J.: The Theory of Affordances. Lawrence Erlbaum (1977)
4. Robotic ubiquitous cognitive network (rubicon) website (2014). http://www.fp7rubicon.eu/
5. Horton, T.E., Chakraborty, A., Amant, R.S.: Affordances for robots: A brief survey. Avant **3**(2), 70–84 (2012)
6. Bacciu, D., Gallicchio, C., Micheli, A., Di Rocco, M., Saffiotti, A.: Learning context-aware mobile robot navigation in home environments. In: The 5th International Conference on Information, Intelligence, Systems and Applications, IISA 2014, pp. 57–62, July 2014