Artificial Intelligence 247 (2017) 45–69

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Artificial cognition for social human–robot interaction: An implementation

Séverin Lemaignan [a,b,*], Mathieu Warnier [a], E. Akin Sisbot [a], Aurélie Clodic [a], Rachid Alami [a]

[a] LAAS-CNRS, Univ. de Toulouse, CNRS, 7 avenue du Colonel Roche, F-31400 Toulouse, France
[b] Centre for Robotics and Neural Systems, Plymouth University, Plymouth, United Kingdom

ARTICLE INFO

Article history:
Received in revised form 4 July 2016
Accepted 9 July 2016
Available online 26 July 2016

Keywords:
Human–robot interaction
Cognitive robotics
Perspective taking
Cognitive architecture
Knowledge representation and reasoning

ABSTRACT

Human–Robot Interaction challenges Artificial Intelligence in many regards: dynamic, partially unknown environments that were not originally designed for robots; a broad variety of situations with rich semantics to understand and interpret; physical interactions with humans that requires fine, low-latency yet socially acceptable control strategies; natural and multi-modal communication which mandates common-sense knowledge and the representation of possibly divergent mental models. This article is an attempt to characterise these challenges and to exhibit a set of key decisional issues that need to be addressed for a cognitive robot to successfully share space and tasks with a human.

We identify first the needed individual and collaborative cognitive skills: geometric reasoning and situation assessment based on perspective-taking and affordance analysis; acquisition and representation of knowledge models for multiple agents (humans and robots, with their specificities); situated, natural and multi-modal dialogue; human-aware task planning; human–robot joint task achievement. The article discusses each of these abilities, presents working implementations, and shows how they combine in a coherent and original deliberative architecture for human–robot interaction. Supported by experimental results, we eventually show how explicit knowledge management, both symbolic and geometric, proves to be instrumental to richer and more natural human–robot interactions by pushing for pervasive, human-level semantics within the robot's deliberative system.

1. The challenge of human–robot interaction

1.1. The human–robot interaction context

Human–Robot Interaction (HRI) represents a challenge for Artificial Intelligence (AI). It lays at the crossroad of many sub-domains of AI and, in effect, it calls for their integration: modelling humans and human cognition; acquiring, representing, manipulating in a tractable way abstract knowledge at the human level; reasoning on this knowledge to make decisions; eventually instantiating those decisions into physical actions both legible to and in coordination with humans. Many AI techniques are mandated, from visual processing to symbolic reasoning, from task planning to *theory of mind* building, from reactive control to action recognition and learning.

* Corresponding author at: Plymouth University, Portland Square A216, PL48AA Plymouth, United Kingdom.
E-mail addresses: severin.lemaignan@laas.fr, severin.lemaignan@plymouth.ac.uk (S. Lemaignan), mathieu.warnier@aas.fr (M. Warnier), e.akin.sisbot@aas.fr (E.A. Sisbot), aurelie.clodic@aas.fr (A. Clodic), rachid.alami@aas.fr (R. Alami).

http://dx.doi.org/10.1016/j.artint.2016.07.002

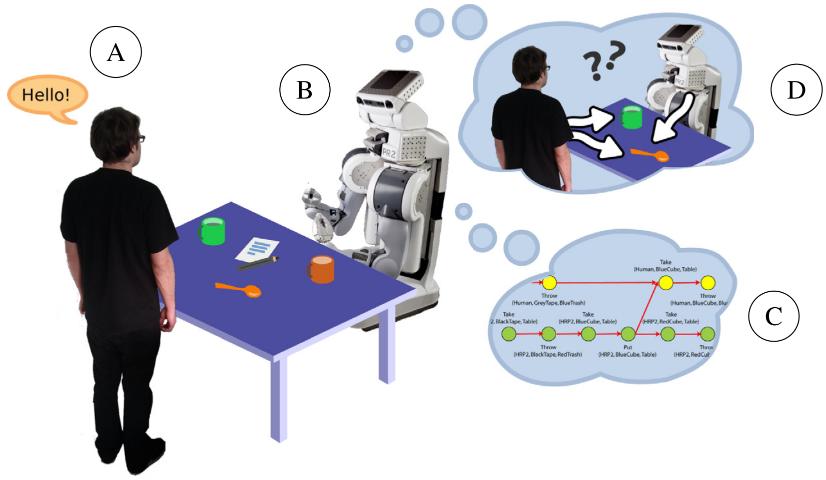

Fig. 1. The robot reasons and acts in domestic interaction scenarios. The sources of information are multi-modal dialogue *(A)* and perspective-aware monitoring of the environment and human activity *(B)*. The robot must adapt on-line its behaviours by merging computed plans *(C)* with reactive control. The robot explicitly reasons on the fact that it is (or is not) observed by the human. Reasoning and planning take place at symbolic as well as geometric level and take into account agents beliefs, perspectives and capabilities *(D)* as estimated by the robot.

We do not claim to address here the issue as a whole. This article attempts however to organise it into a coherent challenge for Artificial Intelligence, and to explain and illustrate some of the paths that we have investigated on our robots, that result in a set of deliberative, knowledge-oriented, software components designed for human–robot interaction.

We focus on a specific class of interactions: human–robot collaborative task achievement [1] supported by multi-modal and situated communication. Fig. 1 illustrates this context: the human and the robot share a common space and exchange information through multiple modalities (we specifically consider verbal communication, deictic gestures and social gaze), and the robot is expected to achieve interactive object manipulation, fetch and carry tasks and other similar chores by taking into account, at every stage, the intentions, beliefs, perspectives, skills of its human partner. Namely, the robot must be able to recognise, understand and participate in communication situations, both explicit (e.g. the human addresses verbally the robot) and implicit (e.g. the human points to an object); the robot must be able to take part in joint actions, both pro-actively (by planning and proposing resulting plans to the human) and reactively; the robot must be able to move and act in a safe, efficient and legible way, taking into account social rules like proxemics.

These three challenges, *communication*, *joint action*, *human-aware execution*, structure the research in human–robot interaction. They can be understood in terms of cognitive skills that they mandate. *Joint action*, for instance, builds from:

- a joint *goal*, which has been previously established and agreed upon (typically through dialogue);
- a physical environment, estimated through the robot's exteroceptive sensing capabilities, and augmented by inferences drawn from previous observations;
- a belief state that includes *a priori* common-sense knowledge and mental models of each of the agents involved (the robot and its human partners).

The robot controller (with the help of a task planner) decides what action to execute next [2], and who should perform it, from the robot or the human (or both in case of a collaborative action such as a handover [3,4]), how it should achieved and what signals should be sensed and/or produced by the robot to facilitate human–robot joint action [5–8]. It finally controls and monitors its execution. The operation continues until the goal is achieved, is declared unachievable or is abandoned by the human [9].

This translates into several decisional, planning, representation skills that need to be available to the robot [10]. It must be able: 1 to represent and manipulate symbolic belief states, 2 to acquire and keep them up-to-date with respect to the state of the world and the task at hand, 3 to build and iteratively refine shared (human–robot) plans, 4 to instantiate and execute the actions it has to perform, and conversely, to monitor those achieved by its human partner.

Besides, such abilities should be designed and implemented in a task-independent manner, and should provide sufficient levels of parametrisation, so that they adapt to various environments, different tasks and variable levels of engagement of the robot, ranging from teammate behaviour to assistant or pro-active helper.

These are the challenges that we will discuss in this article.

1.2. Contribution and article overview

Our main contributions focus on the architecture of the decisional layer of social robots. Specifically, the deliberative architecture of a robot designed to share space and tasks with humans, and to act and interact in a way that supports the human's own actions and decisions. We present hereafter a model of cognitive integration for service robots that: