DAC-h3: A Proactive Robot Cognitive Architecture to Acquire and Express Knowledge About the World and the Self

Clément Moulin-Frier*, Tobias Fischer*, Maxime Petit, Grégoire Pointeau, Jordi-Ysard Puigbo, Ugo Pattacini, Sock Ching Low, Daniel Camilleri, Phuong Nguyen, Matej Hoffmann, Hyung Jin Chang, Martina Zambelli, Anne-Laure Mealiere, Andreas Damianou, Giorgio Metta, Tony J. Prescott, Yiannis Demiris, Peter Ford Dominey, and Paul F. M. J. Verschure

Abstract—This paper introduces a cognitive architecture for a humanoid robot to engage in a proactive, mixed-initiative exploration and manipulation of its environment, where the initiative can originate from both the human and the robot. The framework, based on a biologically-grounded theory of the brain and mind, integrates a reactive interaction engine, a number of state-of-the-art perceptual and motor learning algorithms, as well as planning abilities and an autobiographical memory. The architecture as a whole drives the robot behavior to solve the symbol grounding problem, acquire language capabilities, execute goal-oriented behavior, and express a verbal narrative of its own experience in the world. We validate our approach in human-robot interaction experiments with the iCub humanoid robot, showing that the proposed cognitive architecture can be applied in real time within a realistic scenario and that it can be used with naive users.

Index Terms—Cognitive Robotics, Distributed Adaptive Control, Human-Robot Interaction, Symbol Grounding, Autobiographical Memory

I. INTRODUCTION

THE so-called *Symbol Grounding Problem* (SGP, [1], [2], [3]) refers to the way in which a cognitive agent forms an internal and unified representation of an external word referent

from the continuous flow of low-level sensorimotor data generated by its interaction with the environment. In this paper, we focus on solving the SGP in the context of human-robot interaction (HRI), where a humanoid iCub robot [4] acquires and expresses knowledge about the world by interacting with a human partner. Solving the SGP is of particular relevance in HRI, where a repertoire of shared symbolic units forms the basis of an efficient linguistic communication channel between the robot and the human.

To solve the SGP, several questions should be addressed:

- How are unified symbolic representations of external referents acquired from the multimodal information collected by the agent (e.g., visual, tactile, motor)? This is referred to as the *Physical SGP* [5], [6].
- How to acquire a shared lexicon grounded in the sensorimotor interactions between two (or more) agents? This is referred to as the *Social SGP* [6], [7].
- How is this lexicon then used for communication and collective goal-oriented behavior? This refers to the functional role of physical and social symbol grounding.

This paper addresses these questions by proposing a complete cognitive architecture for HRI and demonstrating its abilities on an iCub robot. Our architecture, called *DAC-h3*, builds upon our previous research projects in conceiving biologically grounded cognitive architectures for humanoid robots based on the Distributed Adaptive Control theory of mind and brain (DAC, presented in the next section). In [8] we proposed an integrated architecture for generating a socially competent humanoid robot, demonstrating that gaze, eye contact and utilitarian emotions play an essential role in the psychological validity or social salience of HRI (DAC-h1). In [9], we introduced a unified robot architecture, an innovative Synthetic Tutor Assistant (STA) embodied in a humanoid robot whose goal is to interactively guide learners in a science-based learning paradigm through rich multimodal interactions (DAC-h2).

DAC-h3 is based on a developmental bootstrapping process where the robot is endowed with an intrinsic motivation to act and relate to the world in interaction with social peers. Levinson [10] refers to this process as the *human interaction engine*: a set of capabilities including looking at objects of interest and interaction partners, pointing to these entities [11], demonstrating curiosity as a desire to acquire knowledge [12]

Manuscript received December 31, 2016; revised July 25, 2017; accepted August 09, 2017. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. FP7-ICT-612139 (What You Say Is What You Did project), as well as the ERC's CDAC project: Role of Consciousness in Adaptive Behavior (ERC-2013-ADG 341196). M. Hoffmann was supported by the Czech Science Foundation under Project GA17-15697Y. P. Nguyen received funding from ERC's H2020 grant agreement No. 642667 (SECURE).

C. Moulin-Frier, J.-Y. Puigbo, S. C. Low, and P. F. M. J. Verschure are with the Laboratory for Synthetic, Perceptive, Emotive and Cognitive Systems, Universitat Pompeu Fabra, 08002 Barcelona, Spain.

T. Fischer, M. Petit, H. J. Chang, M. Zambelli, and Y. Demiris are with the Personal Robotics Laboratory, Department of Electrical and Electronic Engineering, Imperial College London, SW7 2AZ, U.K.

G. Pointeau, A.-L. Mealiere, and P. F. Dominey are with the Robot Cognition Laboratory of the INSERM U846 Stem Cell and Brain Research Institute, Bron 69675, France.

U. Pattacini, P. Nguyen, M. Hoffmann, and G. Metta are with the Italian Institute of Technology, iCub Facility, Via Morego 30, Genova, Italy. M. Hoffmann is also with the Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague.

D. Camilleri, A. Damianou, and T. J. Prescott are with the Department of Psychology, University of Sheffield, U.K.

Digital Object Identifier

*C. Moulin-Frier and T. Fischer contributed equally to this work.

and showing, telling and sharing this knowledge with others [11], [13]. These are also coherent with the *desiderata for developmental cognitive architectures* proposed in [14] stating that a cognitive architecture's value system should manifest both exploratory and social motives, reflecting the psychology of development defended by Piaget [15] and Vygotsky [16].

This interaction engine drives the robot to proactively control its own acquisition and expression of knowledge, favoring the grounding of acquired symbols by learning multimodal representations of *entities* through interaction with a human partner. In *DAC-h3*, an entity refers to an internal or external referent: it can be either an object, an agent, an action, or a body part. In turn, the acquired multimodal and linguistic representations of entities are recruited in goal-oriented behavior and form the basis of a persistent concept of self through the development of an autobiographical memory and the expression of a verbal narrative.

We validate the proposed architecture following a human-robot interaction scenario where the robot has to learn concepts related to its own body and its vicinity in a proactive manner and express those concepts in goal-oriented behavior. We show a complete implementation running in real-time on the iCub humanoid robot. The interaction depends on the internal dynamics of the architecture, the properties of the environment and the behavior of the human. We analyze a typical interaction in detail and provide videos showing the robustness of our system in various environments (<https://github.com/robotology/wysiwyd>). Our results show that the architecture autonomously drives the iCub to acquire a number of concepts about the present entities (objects, humans, and body parts), whilst proactively maintaining the interaction with a human and recruiting those concepts to express more complex goal-oriented behavior. We also run experiments with naive subjects in order to test the effect of the robot's proactivity level on the interaction.

In Section II we position the current contribution with respect to related works in the field and rely on this analysis to emphasize the specificity of our approach. Our main contribution is described in Section III and consists in the proposal and implementation of an embodied and integrated cognitive architecture for the acquisition of multimodal information about external word referents, as well as a context-dependent lexicon shared with a human partner and used in goal-directed behavior and verbal narrative generation. The experimental validation of our approach on an iCub robot is provided in Section IV, followed by a discussion in Section V.

II. RELATED WORKS AND PRINCIPLES OF THE PROPOSED ARCHITECTURE

Designing a cognitive robot that is able to solve the SGP requires a set of heterogeneous challenges to be addressed. First, the robot has to be driven by a cognitive architecture bridging the gap between low-level reactive control and symbolic knowledge processing. Second, it needs to interact with its environment, including social partners, in a way that facilitates the acquisition of symbolic knowledge. Third, it needs to actively maintain engagement with the social partners for a fluent interaction. Finally, the acquired symbols need to be

used in high-level cognitive functions dealing with linguistic communication and autobiographical memory.

In this section, we review related works on each of these topics along with a brief overview of the solution adopted by the *DAC-h3* architecture.

A. Functionally-driven vs. biologically-inspired approaches in social robotics

The methods used to conceive socially interactive robots derive predominantly from two approaches [17]. Functionally-designed approaches are based on reverse engineering methods, assuming that a deep understanding of how the mind operates is not a requirement for conceiving socially competent robots (e.g. [18], [19], [20]), whilst biologically-inspired robots are based on theories of natural and social sciences and expect two main advantages of constraining cognitive models by biological knowledge: to conceive robots that are more understandable to humans, as they reason using similar principles, and to provide an efficient experimental benchmark from which the underlying theories of learning can be confronted, tested and refined (e.g. [21], [22], [23]). One specific approach used by Demiris and colleagues for the mirror neuron system is decomposing computational models implemented on robots into *brain operating principles* which can then be linked and compared to neuroimaging and neurophysiological data [23].

The proposed *DAC-h3* cognitive architecture takes advantage of both methods. It is based on an established biologically-grounded cognitive architecture of the brain and the mind (the *DAC* theory, presented below) that is adapted for the HRI domain. However, while the global structure of the architecture is constrained by biology, the implementation of specific modules can be driven by their functionality, i.e. using state-of-the-art methods from machine learning that are powerful at implementing particular functions without being directly constrained by biological knowledge.

B. Cognitive architectures and the SGP

Another distinction in approaches for conceiving social robots, which is of particular relevance for addressing the SGP, reflects a divergence from the more general field of cognitive architectures (or *unified theories of cognition* [24]). Historically, two opposing approaches have been proposed to formalize how cognitive functions arise in an individual agent from the interaction of interconnected information processing modules in a cognitive architecture. Top-down approaches rely on a symbolic representation of a task, which has to be decomposed recursively into simpler ones to be executed by the agent. These rely principally on methods from symbolic artificial intelligence (from the General Problem Solver [25] to Soar [26] or ACT-R [27]). Although relatively powerful at solving abstract symbolic problems, top-down architectures are not able to solve the SGP *per se* because they presuppose the existence of symbols. Thus they are not suitable for addressing the problem of how these symbols can be acquired from low-level sensorimotor signals. The alternative, bottom-up approaches instead implement behavior without relying on complex knowledge representation and reasoning. This is

typically the case in behavior-based robotics [28], emphasizing lower-level sensory-motor control loops as a starting point of behavioral complexity as in the Subsumption architecture [29]. These approaches are not suitable to solve the SGP either because they do not consider symbolic representation as a necessary component of cognition (referred as *intelligence without representation* in [28]).

Interestingly, this historical distinction between top-down representation based and bottom-up behavior based approaches still holds in the domain of social robotics [30], [31]. Representation based approaches rely on the modeling of psychological aspects of social cognition (e.g. [32]), whereas behavior based approaches emphasize the role of embodiment and reactive control to enable the dynamic coupling of agents [33]. Solving the SGP, both in its physical and social aspects, therefore requires the integration of bottom-up processes for acquiring and grounding symbols in the physical interaction with the (social) environment and top-down processes for taking advantage of the abstraction, reasoning and communication abilities provided by the acquired symbol system. This has been referred to as the *micro-macro loop*, i.e. a bilateral relationship between an emerged symbol system at the macro level and a physical system consisting of communicating and collaborating agents at the micro level [34].

Several contributions in social robotics rely on such hybrid architectures integrating bottom-up and top-down processes (e.g. [35], [36], [37], [38]). In [35], an architecture called embodied theory of mind was developed to link high-level cognitive skills to the low-level perceptual abilities of a humanoid and implementing joint attention and intentional state understanding. In [36], or [37], the architecture combines deliberative planning, reactive control, and motivational drives for controlling robots in interaction with humans.

In this paper, we adopt the principles of the *Distributed Adaptive Control* theory of the mind and the brain (DAC, [39], [40]). DAC is a hybrid architecture which posits that cognition is based on the interaction of four interconnected control loops operating at different levels of abstraction (see Figure 1). The first level is called the *somatic layer* and corresponds to the embodiment of the agent within its environment, with its sensors and actuators as well as the physiological needs (e.g. for exploration or safety). Extending bottom-up approaches with drive reduction mechanisms, complex behavior is bootstrapped in DAC from the self-regulation of an agent's physiological needs when combined with reactive behaviors (the *reactive layer*). This reactive interaction with the environment drives the dynamics of the whole architecture [41], bootstrapping learning processes for solving the physical SGP (the *adaptive layer*) and the acquisition of higher-level cognitive representations such as abstract goal selection, memory and planning (the *contextual layer*). These high-level representations in turn modulate the activity at the lower levels via top-down pathways shaped by behavioral feedback. The control flow in DAC is therefore distributed, both from bottom-up and top-down interactions between layers, as well as from lateral information processing within each layer.

C. Representation learning for solving the SGP

As we have seen, a cognitive architecture solving the SGP needs to bridge the gap between low-level sensorimotor data and symbolic knowledge. Several methods have been proposed for compressing multimodal signals into symbols. A solution based on geometrical structures was offered by Gärdenfors with the notion of *conceptual spaces* (e.g., [42]), whereby similarity between concepts is derived from distances in this space. Lieto et al. [43] advocate the use of the conceptual spaces as the *lingua franca* for different levels of representation.

Another approach has been proposed in [44], [45], which involves a single class of mental representations called “Semantic Pointers”. These representations are particularly suited in solving the SGP as they support binding operations of various modalities, which in turn result in a single representation. This representation (which might have been initially formed by an input of a single modality) can then trigger a corresponding concept, whose occurrence leads to simulated stimuli in the other modalities. Furthermore, while semantic pointers can be represented as vectors, the vector representation can be transformed in neural activity which makes the implementation biologically plausible and allows mapping to different brain areas.

Other approaches consider symbols as fundamentally sensorimotor units. For example, Object-Action Complexes (OACs) build symbolic representations of sensorimotor experience and behaviors through the learning of object affordances [46] (for a review of affordance-based approaches, see [47]). In [48], a framework founded on joint perceptuo-motor representations is proposed, integrating declarative episodic and procedural memory systems for combining experiential knowledge with skillful know-how.

In DAC-h3, visual, tactile, motor and linguistic information about the present entities is collected proactively through reactive control loops triggering sensorimotor exploration in interaction with a human partner. Abstract representations are learned on-line using state-of-the-art machine learning methods in each modality (see Section III). An entity is therefore represented internally in the robot's memory as the association between abstracted multimodal representations and linguistic labels.

D. Interaction paradigms and autonomous exploration

Learning symbolic representations from sensorimotor signals requires an autonomous interaction of a robot with the physical and social world. Several interaction paradigms have been proposed for grounding a lexicon in the physical interaction of a robot with its environment. Since the pioneering paradigm of *language games* proposed in [49], a number of multi-agent models have been proposed showing how particular properties of language can self-organize out of repeated dyadic interactions between agents of a population (e.g. [50], [51]).

In the domain of HRI, significant progress has been made in allowing robots to interact with humans, for example in learning shared plans [52], [53], [54], learning to imitate actions [55], [56], [57], [58], and learning motor skills [59] that can be used for engaging in joint activities. Other contributions

have focused on lexicon acquisition through the transfer of sensorimotor and linguistic information from the interaction between a teacher and a learner through imitation [60] or action [61], [62]. However, in most of these interactions, the human is in charge and the robot is following the human's lead: the choice of which concept to learn is left to the human and the robot must identify it. In this case, the robot must solve the referential indeterminacy problem described by Quine [63], where the robot language learner has to extract the external concept that was referred to by the human speaker. However, acquiring symbols by interacting with other agents is not only a unidirectional process of information transfer between a teacher and learner [10].

Autonomous exploration and proactive behavior solves this problem by allowing robots to take the initiative in exploring their environment [64] and interacting with people [65]. The benefit of these abilities for knowledge acquisition has been demonstrated in several developmental robotics experiments. In [66], it is shown how a combination of social guidance and intrinsic motivation improve the learning of object visual categories in HRI. A similar mechanism is adopted in [67] for learning complex sensorimotor mappings in the context of vocal development. In [68], planning conflicts due to the uncertainty of the detected human's intention are resolved by proactive execution of the corresponding task that optimally reduces the system's uncertainty. In [69], the task is to acquire human-understandable labels for novel objects and learning how to manipulate them. This is realized through a mixed-initiative interaction scenario and it is shown that proactivity improves the predictability and success of human-robot interaction.

A central aspect of the *DAC-h3* architecture is the robot's ability to act proactively in a mixed-initiative scenario. This allows self-monitoring of the robot's own knowledge acquisition process, removing dependence on the human's initiative. Interestingly, proactivity in a sense reverses the referential indeterminacy problem mentioned above by shifting the responsibility of solving ambiguities to the agent who is endowed with the adequate prior knowledge to solve it, i.e., the human, in a HRI context. The robot is now in charge of the concepts it wants to learn, and can use joint attention behaviors to guide the human toward the knowledge it wants to acquire. In the proposed system, this is realized through a set of behavioral control loops, by self-regulating knowledge acquisition, and by proactively requesting missing information about entities from the human partner.

E. Language learning, autobiographical memory and narrative expression

We have just described the main components that a robot requires to solve the SGP: a cognitive architecture able to process both low-level sensorimotor data and high-level symbolic representation, mechanisms for linking these two levels in abstract multimodal representations, as well autonomous behaviors for proactively interacting with the environment. The final challenge to address concerns the use of the acquired symbols for higher-level cognition, including language learning, autobiographical memory and narrative expression.

Several works address the ability of language learning in robotics. The cognitive architecture of iTalk [70] focuses on modeling the emergence of language by learning about the robot's embodiment, learning from others, as well as learning linguistic capability. Cangelosi et al. [71] propose that action, interaction and language should be considered together as they develop in parallel, and one influences the others. Antunes et al. [72] assume that language is already learned, and address the issue that linguistic input typically does not have a one-to-one mapping to actions. They propose to perform reasoning and planning on three different layers (low-level robot perception and action execution, mid-level goal formulation and plan execution, and high-level semantic memory) to interpret the human instructions. Similarly, [73] proposes a system to recognize novel objects using language capabilities in one shot. In these works, language is typically used to understand the human and perform actions, but not necessarily to talk about past events that the robot has experienced.

A number of works investigate the expression of past events by developing narratives based on acquired autobiographical memories [74], [75], [76]. In [75], a user study is presented which suggests that a robot's narrative allows humans to get an insight into long term human-robot interaction from the robot's perspective. The method in [76] takes user preferences into account when referring to past interactions. Similarly to our framework, it is based on the implementation and cooperation between both episodic and semantic memories with a dialog system. However, no learning capabilities (neither language nor knowledge) are introduced by the authors.

In the proposed *DAC-h3* architecture, the acquired lexicon allows the robot to execute action plans for achieving goal-oriented behavior from human speech requests. Relevant information throughout the interaction of the robot with humans is continuously stored in an autobiographical memory used for the generation of a narrative self, i.e., a verbal description of the own robot's history over the long term (able to store and verbally describe interactions from a long time ago, e.g. a period of several months).

In the next section, we describe how the above features are implemented in a coherent cognitive architecture made up of functional YARP [77] modules running in real-time on the iCub robot.

III. THE DAC-H3 COGNITIVE ARCHITECTURE

This section presents the *DAC-h3* architecture in detail, which is an instantiation of the DAC architecture for human-robot interaction. The proposed architecture provides a general framework for designing autonomous robots which act proactively for 1) maintaining social interaction with humans, 2) bootstrapping the association of multimodal knowledge with its environment that further enrich the interaction through goal-oriented action plans, and 3) express a verbal narrative. It allows a principled organization of various *functional modules* into a biologically grounded cognitive architecture.

A. Layer and module overview

In *DAC-h3*, the *somatic layer* consists of an iCub humanoid robot equipped with advanced motor and sensory abilities

for interacting with humans and objects. The *reactive layer* ensures the autonomy of the robot through drive reduction mechanisms implementing proactive behaviors for acquiring and expressing knowledge about the current scene. This allows the bootstrapping of adaptive learning of multimodal representations about entities in the *adaptive layer*. More specifically, the *adaptive layer* learns high-level multimodal representations (visual, tactile, motor and linguistic) for the categorization of entities (objects, agents, actions and body parts) and associates them in unified representations. Those representations form the basis of an episodic memory for goal-oriented behavior through planning in the *contextual layer*, which deals with goal representation and action planning. Within the contextual layer, an autobiographical memory of the robot is formed that can be expressed in the form of a verbal narrative.

The complete *DAC-h3* architecture is shown in Figure 1. It is composed of *structural modules* reflecting the cognitive modules proposed by the DAC theory. Each *structural module* might rely on one or more *functional modules* implementing more specific functionalities (e.g. dealing with motor control, object perception, and scene representation). The complete system described in this section, therefore, integrates several state-of-the-art algorithms for cognitive robotics and integrates them into a structured cognitive architecture grounded in the principles of the DAC theory. In the remainder of this section, we describe each *structural module* layer by layer, as well as their interaction with the *functional modules*, and provide references which provide more detail for the individual modules.

B. Somatic layer

The *somatic layer* corresponds to the physical embodiment of the system. We use the iCub robot, an open source humanoid platform developed for research in cognitive robotics [4]. The iCub is a 104cm tall humanoid robot with 53 degrees of freedom (DOF). The robot is equipped with cameras in its articulated eyes allowing stereo vision, and tactile sensors in the fingertips, palms of the hand, arms and torso. The iCub is augmented with an external RGB-D camera above the robot head for agent detection and skeleton tracking. Finally, an external microphone and speakers are used for speech recognition and synthesis, respectively.

The *somatic layer* also contains the physiological *needs* of the robot that will drive its reactive behaviors, as described in the following section on the *reactive layer*.

C. Reactive layer

Following DAC principles, the *reactive layer* oversees the self-regulation of the internal *drives* of a cognitive agent from the interaction of sensorimotor control loops. The drives aim at self-regulating internal state variables (the *needs* of the *somatic layer*) within their respective homeostatic ranges. In biological terms, such an internal state variable could, for example, reflect the current glucose level in an organism, with the associated homeostatic range defining the minimum and maximum values of that level. A drive for eating would then correspond to a

Figure 1. The *DAC-h3* cognitive architecture (see Section III) is an implementation of the DAC theory of the brain and mind (see Section II-B) adapted for HRI applications. The architecture is organized as a layered control structure with tight coupling within and between layers: the somatic, reactive, adaptive and contextual layers. Across these layers, a columnar organization exists that deals with the processing of states of the world or exteroception (left, red), the self or interoception (middle, blue) and action (right, green). The role of each layer and their interaction is described in Section III. White boxes connected with arrows correspond to *structural modules* implementing the cognitive modules proposed in the DAC theory. Some of these structural modules rely on *functional modules*, indicated by acronyms in the boxes next to the structural modules. Acronyms refer to the following functional modules. SR: Speech Recognizer; PASAR: Prediction, Anticipation, Sensation, Attention and Response; AD: Agent Detector; ARE: Action Rendering Engine; OR: Object Recognition; LRH: Language Reservoir Handler; SSM: Synthetic Sensory Memory; PT: Perspective Taking; SRL: Sensorimotor Representation Learning; KSL: Kinematic Structure Learning; OPC: Object Property Collector; ABM: Autobiographical Memory; NSL: Narrative Structure Learning.

self-regulation mechanism where the agent actively searches for food whenever its glucose level is below the homeostatic minimum and stops eating even if food is present whenever it is above the homeostatic maximum. A drive is therefore defined as the real-time control loop triggering appropriate behaviors whenever the associated internal state variable goes out of its homeostatic range, as a way to self-regulate its value in a dynamic and autonomous way.

In the social robotics context that is considered in this paper, the drives of the robot do not reflect biological needs as above but are rather related to knowledge acquisition and expression in social interaction. At the foundation of this developmental bootstrapping process is the intrinsic motivation to interact and communicate. As described by Levinson [10] (see Introduction), a part of the *human interaction engine* is a set of capabilities that include the motivation to interact and communicate through

universal (language independent) manners; including looking at objects of interest and at the interaction partner, as well as pointing to these objects. These reactive capabilities are built into the *reactive layer* of the architecture forming the core of the *DAC-h3* interaction engine. These interaction primitives allow the *DAC-h3* system and the human to share attention around specific entities (body parts, objects, or agents), and to bootstrap learning mechanisms in the *adaptive layer* that associate visual, tactile, motor and linguistic representations of entities as described in the next section.

Currently, the architecture implements the following two drives: one for knowledge acquisition and one for knowledge expression. However, *DAC-h3* is designed in a way that facilitates the addition of new drives for further advancements (see Section V). First, a *drive for knowledge acquisition* provides the iCub with an intrinsic motivation to acquire new knowledge about the current scene. The internal variable associated with this drive is modulated by the number of entities in the current scene with missing information (e.g. unknown name, or missing property). The self-regulation of this drive is realized by proactively asking the human to provide missing information about entities, for instance, their name via speech, synchronized with gaze and pointing; or asking the human to touch the robot's skin associated with a specific body part.

Second, a *drive for knowledge expression* allows the iCub to proactively express its acquired knowledge by interacting with the human and objects. The internal variable associated with this drive is modulated by the number of entities in the current scene without missing information. The self-regulation is then realized by triggering actions toward the known entities, synchronized with verbal descriptions of those actions (e.g. pointing towards an object or moving a specific body part, while verbally referring to the considered entity).

The implementation of these drives is realized through the three structural modules described below, interacting with each other as well as with the surrounding layers: 1) *sensations*, 2) *allostatic controller*, and 3) *behaviors* (see Figure 1).

1) *Sensations*: The *sensations* module deals with low-level sensing to provide relevant information for meaning extraction in the *adaptive layer*. Specifically, the module detects presence and position of other agents and their body parts (*agent detector* functional module; of interest are the head location for gazing at the partner and the location of the hands to detect pointing actions) based on the input of the RGB-D camera. Similarly, it detects objects based on a texture analysis and extracts their location using the stereo vision capabilities of the iCub [78]. The *prediction, anticipation, sensation, attention and response* functional module (PASAR; [79]) calculates the saliency of agents based on their motion (increased velocity leads to increased saliency), and similarly the saliency for objects is increased if they move or the partner points at them. Finally, the *speech recognition* functional module extracts text from human speech sensed by a microphone using the MicrosoftTM Speech API. The functionalities of the *sensations* module can, therefore, be summarized as dimensionality reduction and saliency computation, and the resulting data are used for bootstrapping knowledge in higher layers of the architecture.

2) *Allostatic Controller*: In many situations, several drives which may conflict with each other can be activated at the same time (in the case of this paper, the drive for knowledge acquisition and the drive for knowledge exploration). Such possible conflicts can be solved through the concept of an *allostatic controller* [80], [81], defined as a set of simple homeostatic control loops and dealing with their scheduling to ensure an efficient global regulation of the internal state variables. The scheduling is decided according to the internal state of the robot and the output of the *sensations* module. The decision of which drive to follow depends on several factors: the distance of each drive level to their homeostatic boundaries, as well as predefined drive priorities (in *DAC-h3*, knowledge acquisition has priority over knowledge expression, which results in a curious personality).

3) *Behaviors*: To regulate the aforementioned drives, the *allostatic controller* is connected to the *behaviors* module, and each drive is linked to corresponding behaviors which are supposed to bring it back into its homeostatic range whenever needed. The positive influence of such a drive regulation mechanism on the acceptance of the HRI by naive users has been demonstrated in previous papers [82], [83].

The drive for knowledge acquisition is regulated by requiring information about entities through coordinated behaviors. Those behaviors depend on the type of the considered entity:

- In the case of an object, the robot produces speech (e.g. "What is this object?") while pointing and gazing at the unknown object.
- In the case of an agent, the robot produces speech (e.g. "Who are you?") while looking at the unknown human.
- In the case of a body part, the robot either asks for the name (e.g. "How do you call this part of my body?") while moving it or, if the name is already known from a previous interaction, asks the human to touch the body part while moving it (e.g., "Can you touch my index while I move it, please?").

The multimodal information collected through these behaviors will be used to form unified representations of entities in the *adaptive layer* (see next section).

The drive for knowledge expression is regulated by executing actions towards known entities, synchronized with speech sentences parameterized by the entities' linguistic labels acquired in the *adaptive layer* (see next section). Motor actions are realized through the *action rendering engine* (ARE [84]) functional module which allows executing complex actions such as push, reach, take, look in a coordinated human-like fashion. Language production abilities are implemented in the form of predefined grammars (for example asking for the name of an object). Semantic words associated to entities are not present at the reactive level, but are provided from the learned association operating in the *adaptive layer*. The *iSpeak* module implements a bridge between the iCub and a voice synthesizer by synchronizing the produced utterance with the lip movements of the iCub to realize a more vivid interaction [85].

D. Adaptive layer

The *adaptive layer* oversees the acquisition of a state space of the agent-environment interaction by binding visual, tactile, motor and linguistic representations of entities. It integrates functional modules for maintaining an internal representation of the current scene, visually categorizing entities, recognizing and sensing body parts, extracting linguistic labels from human speech, and learning associations between multimodal representations. They are grouped in three structural modules described below: *perceptions*, *associations* and *action selection* (see Figure 1).

1) *Perceptions*: The *object recognition* functional module [86] is used to learn the categorization of objects directly from the visual information given by the iCub eyes with resort to the most recent deep convolutional networks. The bounding boxes of the objects found in the *Sensations* module are fed to the learning module for the recognition stage. The output of the system consists of the 2D (in the image plane) and 3D (in the world frame) positions of the identified objects along with the corresponding classification scores as stored in the *objects properties collector* memory (explained below).

There are two functional modules related to language understanding and language production, both integrated within the *language reservoir handler* (LRH). The *comprehension of narrative discourse module* receives a sentence and produces the representation of the corresponding meaning, and can thus transform human speech into meaning. The *module for narrative discourse production* receives a representation of meaning and generates the corresponding sentence (meaning to speech). The meaning is represented in terms of PAOR: *predicate(arguments)*, where *arguments* correspond to thematic roles (*agent, object, recipient*). Both models are implemented as recurrent neuronal networks based on reservoir computing [87], [88], [89].

The *synthetic sensory memory* (SSM) module is currently employed for face recognition and action recognition using a fusion of RGB-D data and object location data as provided from the *sensations* module. In terms of action recognition, SSM has been trained to automatically segment and recognize the following actions: push, pull, lift, drop, wave, and point, while also actively recognizing if the current action is known or unknown. More generally, it has been shown that SSM provides abilities for pattern learning, recall, pattern completion, imagination and association [90]. The SSM module is inspired by the role of the hippocampus by fusing multiple sensory input streams and representing them in a latent feature space [91]. During recall, SSM performs classification of incoming sensory data and returns a label along with an uncertainty measure corresponding to the returned label, which is a use case of the action and face recognition tasks within *DAC-h3*. SSM is also capable of imagining novel inputs or reconstructing previously encountered inputs and sending the corresponding generated sensory data. This allows for the replay of memories as detailed in [92].

2) *Associations*: The *associations* structural module produces unified representations of entities by associating the multimodal categories formed in the *perception* module. Those unified representations are formed in the *objects properties*

collector (OPC), a functional module storing all information associated with a particular entity at the present moment in a proto-language format as detailed in [82]. An entity is defined as a concept which can be manipulated and is thus the basis for emerging knowledge. In *DAC-h3*, each entity has a name associated, which might be unknown if the entity has been discovered but not yet explored. More specifically, higher level entities such as objects, body parts and agents have additional intrinsic properties. For example, an object also has a location and dimensions associated with it. Furthermore, whether the object is currently present is encoded as well, and if so, its saliency value (as computed by the PASAR module described in Section III-C). On the other hand, a body part is an entity which contains a proprioceptive property (i.e. a specific joint), and a tactile information property (i.e. association with tactile sensors). Thus, the OPC allows integrating multiple modalities of one and the same entity to ground the knowledge about the self, other agents, and objects, as well as their relations. Relations can be used to link several instances in an ontological model (see Section III-E1: *Episodic Memory*).

Learning the multimodal associations that form the internal representations of entities relies on the behavior generated by the knowledge acquisition drive operating at the *reactive level* (see previous section). Multimodal information about entities generated by those behaviors is bound together by registering the acquired information in the specific data format used by the OPC. For instance, the *language reservoir handler* module described above deals with speech analysis to extract entity labels from human replies (e.g. “this is a cube”; {P:is, A:this, O:cube, R:∅}). The extracted labels are associated with the acquired multimodal information which depends on the entity type: visual representations generated by the *object recognition* module in case of an object or *agent detector* in case of an agent, as well as motor and touch information in case of a body part.

The associations of representations can also be applied to the developmental robot itself (instead of external entities as above), to acquire motor capabilities or to learn the links between motor joints and skin sensors of its body [93]. Learning self-related representations of the robot’s own body schema is realized by the *sensorimotor representation learning* functional module dedicated to forward model learning [94]. The module receives sensory data collected from the robot’s sensors (e.g. cameras, skin, joint encoders) and allows accurate prediction of the next state given the current state and an action. Importantly, the forward model is learned based on sensory experiences rather than based on known mechanical properties of the robot’s body.

The *kinematic structure learning* functional module [95], [96] estimates an articulated kinematic structure of arbitrary objects (including the robot’s body parts and humans) using visual input videos of the iCub eye cameras. This again is based on sensory experiences rather than known properties of the agents, which is important to autonomously identify the abilities of other agents. Based on the estimated articulated kinematic structures [95], we also allow the iCub to anchor two objects’ kinematic structure joints by observing their movements [96] and formulating the problem of finding corresponding kinematic joint matches between two articulated kinematic structures. This

Figure 2. Examples of the kinematic structure correspondences which the iCub has found. The top figure shows the correspondences between the left and right arm of the iCub, which can be used to infer the body part names of one arm if the corresponding names of the other arm are known. Similarly, the bottom figure shows correspondences between the robot's body and the human's body.

allows the iCub to infer correspondences between its own body parts (its left arm and its right arm), as well as between its own body and the body of the human as retrieved by the *agent detector* [93] (see Figure 2).

Finally, based on these correspondences, the *perspective taking* functional module [97] enables the robot to reason about the state of the world from the partner's perspective. This is important in situations where the views of the robot and the human diverge, for example, due to objects which are hidden to the human but visible to the robot. More importantly, perspective taking is thought to be an essential element for successful cooperation and to ease communication, for example by resolving ambiguities [98]. By mentally aligning the self-perspective with that of the human partner, this module allows algorithms (concerned with the visuospatial perception of the world) to reason as if the input was acquired from an egocentric perspective; which allows to use learning algorithms trained on egocentric data to reason on data acquired from the human's perspective without the need of adapting them.

3) *Action Selection*: The *action selection* module uses the information from *associations* to provide context to the *behaviors* module at the reactive level. This context corresponds to entity names which are provided as parameters to the *behaviors* module, for instance pointing at a specific object or using the object linguistic label in the parameterized grammars defined at the reactive level. This module also deals with the scheduling of action plans from the *contextual layer* according to the current state of the system as explained in the following.

E. Contextual layer

The *contextual layer* deals with higher-level cognitive functions that extend the time horizon of the cognitive agent, such as an episodic memory, goal representation, planning and the formation of a persistent autobiographical memory of the robot interaction with the environment. These functions rely on the unified representations of entities acquired at the *adaptive level*. The *contextual layer* consists of three functional modules that are described below: 1) *episodic memory*, 2) *goals and*

action plans, and 3) *autobiographical memory* used to generate a narrative structure.

1) *Episodic Memory*: The *episodic memory* relies on advanced functions of the *object property collector* (OPC) to store and associate information about entities in a uniform format based on the interrogative words “who” (is acting), “what” (they are doing), “where” (it happens), “when” (it happens), “why” (it is happening) and “how” called an H5W data structure [82]. It is used for goal representation and as elements of the *autobiographical memory*. H5W have been argued to be the main questions any conscious being must answer to survive in the world [99], [100].

The concept of *relations* is the core of the H5W framework. It links up to five concepts and assigns them with semantic roles to form a solution to the H5W problem. We define a *relation* as a set of five edges connecting those nodes in a directed and labeled manner. The labels of those edges are chosen so that the *relation* models a typical sentence from the English grammar of the form: Relation $\rightarrow$ Subject Verb [Object] [Place] [Time]. The brackets indicate that the components are optional; the minimal relation is therefore composed of two entities representing a subject and a verb.

2) *Goals and action plans*: Goals can be provided to the iCub from human speech, and a meaning is extracted by the *language reservoir handler*, forming the representation of a goal in the *goals* module. Each goal consequently refers to the appropriately predefined action plan, defined as a state transition graph with states represented by nodes and actions represented by edges of the graph. The *action plans* module extracts sequences of actions from this graph, with each action being associated with a pre- and a post-condition state. Goals and action plans can be parameterized by the name of a considered entity. For example, if the human asks the iCub to take the cube, this loads an action plan for the goal “Take an object” which consists of two actions: “Ask the human to bring the object closer” and “Pull the object”. In this case, each action is associated with a pre- and post-condition state in the form of a region in the space where the object is located. In the *action selection* module of the adaptive layer, the plan is instantiated toward a specific object according to the knowledge retrieved from the *associations* module (e.g. allowing to retrieve the current position of the cube). The minimal sequence of actions achieving the goal is then executed according to the perceived current state updated in real-time, repeating each action until its post-condition is met (or cease making the effort after a predefined timeout).

Although quite rigid in its current implementation, in the sense that action plans are predefined instead of being learned from the interaction, this planning ability allows closing the loop of the whole architecture, where drive regulation mechanisms at the *reactive layer* can now be bypassed through contextual goal-oriented behavior. Limitations of this system are discussed in Section V: Conclusions.

3) *Autobiographical Memory*: The *autobiographical memory* (ABM [101], [102], [103]) collects long term information (days, months, years) about interactions motivated by the human declarative long term memory situated in the medial temporal lobe, and the distinction between facts and events

[101]. It stores data (e.g. objects locations, human presence) from the beginning to the end of an episode by taking snapshots of the environmental information from the *episodic memory* containing the pre-conditions and effects of episodes. This allows the generation of high-level concepts extracted by knowledge-based reasoning. In addition, the ABM captures continuous information during an episode (e.g. images from the camera, joints values) [102], which can be used by reasoning modules focusing on the action itself, leading to the production of a procedural memory (e.g. through learning from motor babbling or imitation) [103].

The *narrative structure learning module* builds on the language processing and ABM capabilities. Narrative structure learning occurs in three phases: 1) First the iCub acquires experience in a given scenario, which generates the meaning representation in the ABM. 2) The iCub then formats each story in term of initial states, goal states, actions and results (IGARF graph [89]). 3) The human then provides a narration (that is understood using the reservoir system explained in Section III-D) for the scenario.

By mapping the events of the narration to the event of the story, the robot can extract the meaning of different discourse functions words (such as “because”). It can thus automatically generate the corresponding form-meaning mapping that defines the individual grammatical constructions and their sequencing that defines the narrative construction of a new narrative.

F. Summary on the DAC-h3 architecture

The DAC-h3 architecture, therefore, integrates several state-of-the-art algorithms for cognitive robotics and integrates them into a structured cognitive architecture grounded in the principles of the DAC theory. The drive reduction mechanisms in the *reactive layer* allow a complex control of the iCub robot which proactively interacts with humans. In turn, this allows the bootstrapping of adaptive learning of multimodal representations about entities in the *adaptive layer*. Those representations form the basis of an episodic memory for goal-oriented behavior through planning in the *contextual layer*. The life-long interaction of the robot with humans continuously feed an autobiographical memory able to retrieve past experience from request and to express it verbally in a narrative. Altogether, this allows the iCub to interact with humans in complex scenarios, as described in the next section.

IV. EXPERIMENTAL RESULTS

This section validates the cognitive architecture described in the previous section on a real demonstration with an iCub humanoid robot interacting with objects and a human. We first describe the experimental setup, followed by the behaviors provided to the robot (self-generated and human-requested). Finally, we analyze the DAC-h3 system in two ways: a complete version reporting the full complexity of the system through multiple video demonstrations and the detailed analysis of a particular interaction, as well as a simplified version showing the effect of the robot’s proactivity level on naive users.

The code for reproducing these experiments on any iCub robot is available open-source at <https://github.com/robotology/>

Figure 3. The setup consists of an iCub robot interacting with objects on the table and a human in front of it. The table is separated (indicated by horizontal lines) into three areas: *I* for the area only reachable by the iCub, *S* for the shared area, and *H* for the human-only area (compare with Section IV-A).

wysiwyd. It consists of all modules described in the last section implemented in either C++ or Python, and relies on the YARP middleware [77] for defining their connections and ensuring their parallel execution in real-time.

A. Experimental setup

We consider an HRI scenario where the iCub and a human face each other with a table in the middle and objects placed on it. The surface of the table is divided into three distinct areas, as shown in Figure 3:

- 1) an area which is only reachable by the iCub (*I*),
- 2) an area which is only reachable by the human (*H*), and
- 3) an area which is reachable by both agents (*S* for *Shared*).

The behaviors available to the iCub are the following:

- “Acquire missing information about an entity”, which is described in more detail in Section IV-B1.
- “Express the acquired knowledge”, which is described in more detail in Section IV-B2.
- “Move an object on the table”, either by pushing it from region *I* to *S* or pulling it from region *S* to *I*,
- “Ask the human to move an object”, either by asking to push the object from region *H* to *S* or by asking to pull it from region *S* to *H*.
- “Show learned representations on screen” while explaining what is being shown, e.g. displaying the robot kinematic structure learned from a previous arm babbling phase.

- “Interact verbally with the human” while looking at her/him. This is used for replying to some human requests as described in Section IV-C.

These behaviors are implemented in the *behaviors* module and can be triggered from two distinct pathways as shown in Figure 1. The behaviors for acquiring and expressing knowledge are triggered through the drive reduction mechanisms implemented in the *allostatic controller* (Section III-C) and are self-generated by the robot. The remaining behaviors are triggered from the *action selection* module (Section III-D), scheduling action sequences from the *goals* and *action plans* modules (Section III-E). In the context of the experiments described in this section, these behaviors are requested by the human partner. We describe these two pathways in the two following subsections.

B. Self-generated behavior

Two drives for knowledge acquisition and knowledge expression implement the interaction engine of the robot (see Section III-C). They regulate the knowledge acquisition process of the iCub and proactively maintain the interaction with the human. The generated sensorimotor data feeds the adaptive layer of the cognitive architecture to acquire multimodal information about the present entities (see Section III-D). In the current experiment, the entities are objects on the table, body parts (fingers of the iCub), human partners, and actions. The acquired multimodal information depends on the considered entity. Object representations are based on visual categorization and stereo-vision based 3D localization performed by the *object recognition* functional module. Body part representations associate motor and touch events. Agents and actions representations are learned from visual input in the *synthetic sensory memory* module presented in Section III-D1. Each entity is also associated with a linguistic label learned by self-regulating the two drives detailed below.

1) *Drive to acquire knowledge*: This drive maintains a curiosity-driven exploration of the environment by proactively requesting the human to provide information about the present entities, e.g. naming an object or touching a body part. The drive level decays proportionally to the amount of missing information about the present entities (e.g. the unknown name of an entity). When below a given threshold, it triggers a behavior following a generic pattern of interaction, instantiated according to the nature of the knowledge to be acquired. It begins with a behavior to obtain a joint attention between the human and the robot toward the entity that the robot wants to learn about. After the attention has been attracted toward the desired entity, the iCub asks for the missing information (e.g. the name of an object or of the human, or in the case of a body part the name and touch information) and the human replies accordingly. In a third step, this information is passed to the adaptive layer and the knowledge of the robot is updated in consequence.

Each time the drive level reaches the threshold, an entity is chosen in a pseudo-random way within the set of perceived entities with missing information, with a priority to request the name of a detected unknown human partner. Once a new agent

enters the scene, the iCub asks for her/his name, which is stored alongside representations of its face in the *synthetic sensory memory* module. Similarly, the robot stores all objects it has previously encountered in its *episodic memory* implemented by the *object property collector* module. When the chosen entity is an object, the robot asks the human to provide the name of interest while pointing at it. Then, the visual representation of the object computed by the *object recognition* module is mapped to the name. When the chosen entity is a body part (left-hand fingers), the iCub first raises its hand and moves a random finger to attract the attention of the human. Then it asks for the name of that body part. This provides a mapping between the robot’s joint identifier and the joint’s name. This mapping can be extended to include tactile information by asking the human to touch the body part which is being moved by the robot.

Once a behavior has been triggered, the drive is reset to its default value and decays again as explained above (the amount of the decay being reduced according to what has been acquired).

2) *Drive to express knowledge*: This drive regulates how the iCub expresses the acquired knowledge through synchronized speech, pointing and gaze. It aims at maintaining the interaction with the human by proactively informing her/him about its current state of knowledge. The drive level decays proportionally to the amount of already acquired information about the present entities. When below a given threshold (meaning that a significant amount of information has been acquired), it triggers a behavior alternating gazing toward the human and a known entity, synchronized with speech expressing the knowledge verbally, e.g. “This is the octopus”, or “I know you, you are Daniel”. Once such a behavior has been triggered, the drive is reset to its default value and decays again as explained above (the amount of the decay changing according to what is learned by satisfying the drive for knowledge acquisition).

These two drives allow the robot to balance knowledge acquisition and expression in an autonomous and dynamic way. At the beginning of the interaction, the robot has little knowledge about the current entities and therefore favors behaviors for knowledge acquisition. By acquiring more and more knowledge, it progressively switches to behaviors for knowledge expression. If new entities are introduced, e.g. a new object or another human, it will switch back to triggering more behaviors for knowledge acquisition and so on.

C. Human-requested behavior

The representations which are acquired through satisfying the drives introduced above allow a more complex interaction through goal-oriented behavior managed by the contextual layer (see Figure 1 and Section III-E). Goals can be provided to the iCub from human speech and a meaning is extracted by the *language reservoir handler*, forming the representation of a goal in the *goals* module. Each goal is associated with an action plan on the form of a sequence of actions together with their pre- and post-conditions in the *action plans* module. The *action selection* module takes care of the execution of the plan according to the *associations* known to the robots, triggering

Figure 4. State transition graph used for generating the action plans of the goals “Give” and “Take”. Nodes correspond to the table regions indicated in Figure 3. Arrows correspond to the actions to be executed for realizing a transition. According to the observed current state of the object, the iCub will execute the action which brings it closer to the goal state. For example, if the goal is to take an object which is in the human area, the robot will first ask the human to push it, and subsequently pull it into its own area.

the appropriate behaviors according to its current perception of the scene updated in real time. Goal achievement bypasses the reactive behavior described in the previous subsection by freezing all the drive levels during the execution of the plan. The available goals are described below.

1) *Give or take an object*: These goals are generated from a human verbal request, e.g. “Give me the octopus” or “Take the cube”. Here, the goal is represented as a region on the table, either the human area *H* (for the “Give” goal) or the iCub area *I* (for the “Take” goal), where the mentioned object should be placed. Action plans are generated from the state-transition graph shown in Figure 4. State perception is updated in real-time according to the perceived location of the object computed through stereo-vision in the *object recognition* module.

2) *Point to an object*: This goal is generated through a verbal request, e.g. “Point to the octopus”. If the mentioned object is not known to the iCub, it will first ask the human to point to it to learn the new association between the name and the object’s visual representation. Once the name is known, or if it was already known, the iCub will point to the object.

3) *Say the name of a recognized action*: This goal is generated through a verbal request, e.g. “How do you call this action?” formulated just after the human has performed an action on an object. Six actions can be recognized by the *synthetic sensory memory* module: “push”, “pull”, “lift”, “drop”, “wave”, and “point”. The reply from the iCub provides the name of the action and the object as well as the hand used by the human, e.g. “You pushed the cube with your left hand”.

4) *Say what happened during a past interaction with a human*: This goal is generated through a verbal request, e.g. “What have you done the other day?”. Based on its past interactions with the environment and with humans, the iCub has stored all the relevant information in its *autobiographical memory* (see Section III-E), including its own drives, motivations, and actions, as well of the actions of the human (both, spoken and physically enacted). The *narrative handler* functional module can generate a narrative discourse from the content of the *autobiographical memory* and to generate an action plan in the form of a sequence of sentences. The human can request more detailed information about an event using sentences like “What happened next?” or “Why did you do that?”, this later question being answered according to the stored drive levels and goal states of the considered events. The robot can also learn a set of questions that it can re-use in the context of another story. Figure 5 shows an example of a generated narrative.

First **I** wanted to get the toy.
 First **you** have the toy.
 Then **I** fail to grasp it.
 After **I** fail to grasp, **I** reasoned.
 Because **I** reasoned, **I** ask for the toy to you
 Finally **you** gave me the toy
 Now **I** have the toy
 I have the toy because **you** gave it to me
 You gave the toy to me because **I** wanted it

Figure 5. Example of a narrative generated by the robot. The *language reservoir handler* will decompose the words in the narrative discourse in 3 categories: the *discourse function words* (DFW) which direct the discourse from one sentence to the other, the *open class words* (OCW) which correspond to the meaningful words in terms of vocabulary of the sentence, and the *closed class words* (CCW) which have a grammatical function in the sentence (see [89]).

5) *Show the learned kinematic structure*: As for the previous goals, this goal is generated through verbal requests. When asked “What have you learned from your arm babbling?”, the iCub refers the human to look at the screen where the kinematic structures of its arms are displayed. Also, lines which connect nodes of the kinematic structures indicate the correspondences which the iCub has found between its left and right arm. Similarly, the iCub displays the correspondences which it has found between one of its arms and the body of the human (see Figure 2). This knowledge is further employed to point to the human’s arm, which is interesting as both the name as well as the kinematic location of the human’s arm are inferred from self-learned representations and mapping these representations to the partner.

D. Scenario Progression

We first show how the full DAC-h3 system we have just described is able to generate a complex human-robot interaction by providing videos of live interactions (see <https://github.com/robotology/wysiwyd>) and a detailed description of a particular interaction. Then, in the next subsection, we will analyze more specifically the effect of the robot’s proactivity level on naive users. In both cases, we consider a mixed-initiative scenario, where the iCub behaves autonomously as described in Section IV-B, and so does the human. The human can interrupt the robot behavior by formulating verbal requests as described in Section IV-C. The scenario can follow various paths according to the interaction between the iCub’s internal drive dynamics, its perception of the environment, and the behavior of the human.

Here, we describe one particular instance of the scenario. Figure 6 shows the corresponding drive dynamics and human-robot interactions, and Figure 1 shows the connections between the modules of the cognitive architecture. Each of the numbered items below refers to its corresponding number in Figure 6.

- 1) At the beginning of the interaction, the iCub has only limited knowledge about the current scene. In the *sensations* module, the *agent detector* detects the presence of a human and extracts its skeleton. The *object recognition* module performs blob detection for extracting objects on the table from the visual input of the eye cameras and

Figure 6. Drive level dynamics during a typical mixed-initiative scenario described in Section IV-D. Each drive starts at its default value and decays following the dynamics described in Section IV-B. When reaching a given threshold (dashed horizontal line) the associated behavior is triggered (green rectangles), the corresponding drive level is reset to its default value and both drive levels are frozen for the duration of the behavior. Human behavior is indicated by the red rectangles, being either a reply to a question asked by the iCub (small rectangles) or a request to the iCub triggering goal oriented behavior (here: “Take the cube”). The numbers refer to the description of the scenario progression in the main text.

trains a classifier to categorize them. The 3D locations of the present objects are also computed in *object recognition* through stereo vision. This forms a first incomplete representation of the scene in the *episodic memory* where the *object property collector* registers the location and type of each detected entity (here objects and an agent). It also contains slots for unknown body parts of the iCub, here the five fingers of its right hand.

- 2) The presence of a large amount of missing information (presence of unknown objects, human and body parts) in the *sensations* module makes the drive for knowledge acquisition to decay rapidly in the *allostatic controller* and the drive for knowledge expression is kept constant (since there is no knowledge to express yet).
- 3) When the knowledge acquisition drive level is below threshold, it triggers the associated behavior (*behaviors* module) for acquiring information about an entity. The choice of the unknown entity is pseudo-random, with priority for requesting the name of an unknown human. This makes the robot look at the human. The visual input is passed to the *perceptions* module where the *synthetic sensory memory* segments the face from the background and attempts to recognize it from previously seen faces. If it does not recognize the face, the robot asks, “I do not know you, who are you?”. The human can then reply, e.g., saying “I am Daniel”. The level of the drive is reset to its default value and both drives are frozen during the behavior execution.
- 4) The perceived speech is analyzed by the *language reservoir handler* in *perceptions* to extract the name “Daniel” and is associated with the face representation in the *associations* module. Thus, the next time the iCub will interact with this person, it will directly recognize him and not ask for his name.
- 5) Once this interaction is achieved, the drives continue to decay. Since the iCub has just acquired more information, the decay of the drive for knowledge acquisition is slower

and the one for knowledge expression is increased. Still, the drive for knowledge acquisition reaches the threshold first. The behavior for acquiring information is therefore triggered again. This time, the random choice of an unknown entity makes the robot point to an object and ask, “What is this object?”. The human replies e.g. “This is the cube”. The *language reservoir handler* extracts the name of the object from the reply and the *associations* module associates it with the visual representation of the pointed object from *object recognition*. Now the cube can later be referred by its name.

- 6) The drives continue to decay. This time, the drive for knowledge expression reaches the threshold first. This triggers the behavior for expressing the acquired knowledge. A known entity is chosen, in this example it is the cube, which the robot points at while saying “This is a cube”.
- 7) The human asks “Take the cube”. A meaning is extracted by the *language reservoir handler* in *perceptions* and forms the representation of a goal to achieve in the *goal* module (here the desired location of the object, i.e. the region of the iCub *I* for the goal “take”, see Figure 4). An action plan is built in *action plans* with the sequence of two actions “Ask the human to push the object” then “Pull the object”, together with their pre- and post-conditions in term of states (*I*, *S* or *H*). The *action selection* module takes care of the realization of the plan. First, it instantiates the action plan toward the considered object, here the cube, through its connection with *associations*. Then, it executes each action until its associated post-condition is met (repeating it up to three times before giving up). Since the cube is in the human area *H*, the iCub first triggers the behavior for interacting verbally with the human, asking “Can you please bring the cube closer to the shared area?”. The human pushes the cube to the shared area *S* and the state transition is noticed by the robot thanks to the real-time object localization performed in the *object recognition* module. Then the robot triggers a motor action to pull the cube. Once the goal is achieved (i.e. the cube is in *I*), the drive levels which were frozen during this interaction continue to decay.
- 8) The drive for knowledge acquisition reaches the threshold first. The associated behavior now chooses to acquire the name of a body part. The robot triggers the behavior for raising its hand and moving a random unknown body part, here the middle finger. It looks at the human and asks “How do you call this part of my body?”. The name of the body part is extracted from the human’s reply and is associated with the joint that was moved in *associations*.

The interaction continues following the dynamics of the drives and interrupted by the requests from the human. Once all available information about the present entities is acquired, the drive for knowledge acquisition stops to decay. However, the robot still maintains the interaction through its drive for knowledge expression and the human can still formulate requests for goal-oriented behavior. When new entities are introduced, e.g. an unknown object or another human entering

the scene, the drive for knowledge acquisition decays again and the process continues.

E. Effect of the robot's proactivity level on naive users

We now test the DAC-h3 architecture with naive users having to perform a collaborative task with the iCub robot. For this aim, we conduct experiments on six subjects and compare different configurations of the DAC-h3 drive system reflecting different levels of robot's proactivity.

To better control the experiment, we simplify the setup described above by limiting it to object tagging and manipulation. This means that in this study we do not use the functionalities for agent or body part tagging, action recognition, kinematic structure learning and narrative discourse. The iCub can only proactively tag or point at objects on the table, whereas the human can reply to object tagging requests and provide orders for taking or giving an object. These orders trigger action plans combining object manipulation by the iCub with requests to the human to do so, as described above. We note that, due to the distributed implementation of the DAC-h3 systems as interacting YARP modules operating in parallel, deactivating functionalities simply requires to not launch the corresponding modules and does not imply any modification in the code.

Three objects are placed on the table: an *octopus*, a *blue cube* and a *duck*. Initially, the names of the objects are unknown to the iCub and they are placed on the table as shown on Figure 7 (left). The task given to the subjects is to achieve a goal configuration of object positions (Figure 7, right). The experiment terminates when the task is achieved.

To do so, the subject is instructed to interact with the iCub in the following way. At any time of the interaction, s/he can provide speech orders on the form of the sentences "Give me the <object name>" and "Please take the <object name>", where <object name> is the name of one of the three objects. These names are provided to the subject before the experiment starts to make the speech recognition more robust. Whenever the iCub asks the subject for the name of an object, s/he can reply with "This is the <object name>". Moreover, whenever the iCub asks to show a specific object, the subject can point to it using her/his right hand. This latter behavior is added to the state transition graph described in Figure 4 and executed when the subject requests to perform an action on an object which is unknown to the iCub. To augment the difficulty of the task, the subject is asked to not move objects on its own initiative, but only when the iCub asks to do so.

Achieving the goal configuration requires a complex interaction between the subject and iCub. For example, moving the octopus from the human region to the iCub region requires to first inform the iCub about which object the octopus is, and then asking the iCub to take that object. Informing the robot about the name of an object can occur either from the iCub's initiative through the knowledge acquisition drive, or from the human's initiative by requesting an action on this object (when the object is unknown to the iCub, it will first ask the human to point at it). Since the octopus is not within the reach of the iCub, the robot will first ask the human to move it to the shared region, before executing the motor action for pulling the object in the iCub region.

Figure 7. The experimental task. Each panel shows a top view of the table, where the three letters *I*, *S* and *H* indicates the iCub, Shared and Human areas as in Figure 3. Starting from the initial configuration of object positions on the left, the subject is asked to achieve the goal configuration on the right. This is done by interacting with the iCub following the instructions described above. The three objects are: an octopus (orange object on the left of the table), a blue cube (blue object in the middle) and a duck (yellow object on the right).

We run the experiment with six naive subjects on three different conditions. The three conditions correspond to three different levels of the robot's proactivity, defined by setting different drive decays in the *allostatic controller* module. The two drives for knowledge acquisition and knowledge expression are initialized to a default value of 0.5. Then, they decay linearly at a rate of $n_{obj} * \delta$ unit/s, where δ is a constant value defining the proactivity level. For the knowledge acquisition drive, n_{obj} is the number of perceived objects which are unknown to the robot. For the knowledge expression drive, it is the number of known objects. Therefore, the drive for knowledge acquisition (resp. knowledge expression) decays proportionally to the number of unknown objects (resp. known objects). We define three conditions: medium proactivity (δ for knowledge acquisition = 0.01, for knowledge expression = 0.004), slow proactivity (δ for both drives are 2.5 times lower than in medium proactivity) and fast proactivity (δ for both drives are 2.5 times higher than in medium proactivity). Corresponding behaviors are triggered when a drive value goes below 0.25. For example, for the knowledge acquisition drive in the medium-proactive condition ($\delta = 0.01$) at the beginning of the interaction when all objects are unknown ($n_{obj} = 3$), it takes approximately 8 seconds for the drive to decay completely from the default value of 0.5 to the threshold 0.25.

Figure 8 shows the interaction diagrams for the six subjects grouped by condition. In the slow-proactive condition, we observe that the task can be solved very rapidly, as seen by the subject *S1*. Here the iCub has initiated only the first action (*Tagging (knowledge acquisition)*). With drives decaying slower, the robot acts less often by its own initiative, which leads to the subject leading the interaction according to its own goal. Since the iCub is rarely acquiring information by itself, the subject has to request actions on objects which are not yet known to the robot, triggering tagging behaviors (*Tagging (human points)*) prior to asking or moving objects. However, solving the task can also be quite long in this condition as seen by subject *S2*. This can be due to several factors such as the robot's failure to perceive or manipulate objects, or the time required by the subject to fully apprehend the system. This can also be due to the personality or the mood of the subject who can her/himself be more or less proactive.

At the other extreme, in the fast proactive condition (bottom of the figure), we observe that the interaction is dominated

Figure 8. Interaction diagrams for the six subjects ($S1$ to $S6$, two per condition) interacting with the iCub to solve the task described in Figure 7. Top: slow proactivity condition. Middle: medium proactivity. Bottom: fast proactivity. In each condition, the first five rows shows the interaction of one subject (e.g. $S3$ in the medium proactivity condition) and the five last rows the interaction of another subject (e.g. $S4$ in the medium proactivity condition). For each subject, the colored bars show the times (x-axis) when specific actions (y-axis) are executed by the iCub. The two first rows (*Tagging (knowledge acquisition)* and *Pointing (knowledge expression)*) are actions initiated by the iCub through its drive regulation system in the allostatic controller as described in section IV-B. The three last rows (*Tagging (human points)*, *Robot moves object* and *Ask human to move object*) are actions initiated by the human through a speech order and executed by the iCub through action plans in the contextual layer, as described in section IV-C. They are sequenced according to the current state of the world and the current knowledge of the robot. The vertical arrow in each subject's interaction plot shows the time at which all the three object names are known to the robot.

by the iCub's self-generated actions (*Tagging (knowledge acquisition)* and *Pointing (knowledge expression)*). There is little time for the subject to request orders through speech between two consecutive actions of the iCub and it requires some time to succeed doing it. In consequence, solving the task can require quite a long time as seen by the subject $S6$. However, the names of all objects tend to be acquired faster in this condition (vertical arrows). This is because the drive for knowledge acquisition decays rapidly, especially at the beginning of the interaction where all objects are unknown, pushing the iCub to ask the name of everything around. Indeed, the three first actions are *Tagging (knowledge acquisition)* for both $S5$ and $S6$.

Finally, in the medium-proactive condition, we observe a

mixed initiative between the iCub and the subject. A possible positive effect of such a moderate proactivity of the robot is to maintain the social interaction when the subject is less proactive by her/himself, preventing long silent gaps in the interaction but still largely allowing the human to take the initiative. We observe a better interplay between iCub- and human-initiated actions in this condition.

Subjects are also asked to fill pre- and post-questionnaires before and after the interaction. The pre-questionnaire attempts at characterizing the personality traits of the subject, while the post-questionnaire evaluates subjective feelings about the interaction. As the main focus of this paper is on presenting a coherent cognitive architecture rather than the evaluation of the emerging human-robot interaction, we do not analyze these

data in this paper. Providing a solid statistical analysis will require acquiring data of many more subjects which will be the aim of a follow-up paper. However, the presented human-robot interaction study shows that proactivity is an important factor in HRI and that it can be controlled by simple modifications of the drive dynamics. As observed in Figure 8, the level of robot's proactivity has a clear influence on the resulting interaction by modulating the interplay between the robot and human initiatives.

V. CONCLUSION AND FUTURE WORKS

This paper has introduced *DAC-h3*, a proactive robot cognitive architecture to acquire and express knowledge about the world and the self. The architecture is based on the Distributed Adaptive Control (DAC) theory of the brain and mind, which provides a biologically grounded framework for organizing various functional modules into a coherent cognitive architecture. Those modules implement state-of-the-art algorithms modeling various cognitive functions for autonomous self-regulation, whole-body motor control, multimodal perception, knowledge representation, natural language processing, and goal-oriented behavior. They are all implemented using the YARP robotic middleware on the iCub robot, ensuring their parallel execution in real time and providing synchronous and asynchronous communication protocols among modules.

The implementation of DAC-h3 is flexible so that existing modules can easily be replaced by more sophisticated ones in the future. Moreover, most modules can be configured according to the user's needs, for example by adding new drives or more complex grammars into the system. This makes DAC-h3 a general framework for designing autonomous robots, especially in HRI setups. The underlying open-source code contains some modules which are specific to the iCub robot (modules related to action execution; which work on both real and simulated iCub's), but all cognitive nodes (autobiographical memory, allostatic controllers, etc.) can be freely used with other robots (most easily with YARP-driven robots, but thanks to YARP-ROS intercommunication also with ROS driven robots).

The main contribution of this paper is not about the modeling of the specific functional modules, which already have been published (see Section III), but rather about the integration of a heterogeneous collection of modules into a coherent and operational cognitive architecture. For this aim, the *DAC-h3* architecture is organized as a layered control structure with tight coupling within and between layers (Figure 1 and Section III): the *somatic*, *reactive*, *adaptive*, and *contextual* layers. Across these layers, a columnar organization exists that deals with the processing of states of the world or exteroception, the self or interoception, and action. Two main control loops generate the behavior of the robot. First, a reactive-adaptive control loop ensures autonomy and proactivity through the self-regulation of internal drives for knowledge acquisition and expression. It allows the robot to proactively manage its own knowledge acquisition process and to maintain the interaction with a human partner, while associating multimodal information about entities with their linguistic labels. Second,

an adaptive-contextual control loop allows the robot to satisfy human requests, triggering goal-oriented behavior relying on the acquired knowledge. Those goal-oriented behaviors are related to action planning for object passing, pointing, action recognition, narrative expression and kinematic structure learning demonstration.

We have shown that these two control loops lead to a well-defined interplay between robot-initiated and human-initiated behaviors, which allows the robot to acquire multimodal representations of entities and link them with linguistic symbols, as well as to use the acquired knowledge for goal-oriented behavior. This allows the robot to learn reactively as well as proactively. Reactive learning occurs in situations where the robot requires obtaining new knowledge to execute a human order (e.g. grasping an object with an unknown label), and thus leads to an efficient interaction for acquiring the information before acting according to the human desire. At the same time, the robot can also learn proactively to optimize its cognitive development by triggering learning interactions itself, which allows the robot to learn without having to wait for the human to teach new concepts. Moreover, this is supposed to reduce the cognitive load of the human teacher, as the robot 1) chooses the target entity of the learning interaction, 2) engages the joint attention by making the human aware of the target entity, and 3) asks for the corresponding label. Therefore, the human only needs to provide the label, without having to be concerned about the prior knowledge of the robot.

We have implemented the entire *DAC-h3* architecture and presented an HRI scenario where an iCub humanoid robot interacts with objects and a human to acquire information about the present objects and agents as well as its own body parts. We have analyzed a typical interaction in detail, showing how *DAC-h3* is able to dynamically balance the knowledge acquisition and expression processes according to the properties of the environment, and to deal with a mixed initiative scenario where both the robot and the human are behaving autonomously. In a series of video recordings, we show the ability of *DAC-h3* to adapt to different situations and environments. We have also conducted experiments with naive subjects on a simplified version of the scenario, showing how the robot's proactivity level influences the interaction.

Adapting the proactivity level also provides a step towards personalities of robots. The curiousness (i.e. favoring proactive learning) or talkativeness (i.e. communicating about its own knowledge) of the robot is determined by the decay rates of the corresponding drives. Thus the personality of the robot can be altered by a simple modification of the decay values, as done for skill refinement by Puigbo et al. [104].

The current work has the following limitations. First, some of the available abilities deserve to be better integrated into the HRI scenario. For example, this is the case for the kinematic structure learning process which is currently executed in a separated learning phase instead of being fully integrated within the interaction scenario. Similarly, the narrative can only be generated from specific chunks of the robot's history as recorded in the autobiographical memory. Second, in this paper, we do not provide a statistical analysis of the HRI experiments. The reason is that we focus on the description

of the entire architecture and on their theoretical principles. A thorough statistical analysis will require collecting much more data to fully demonstrate the utility of some of these principles, for example how proactivity helps to solve the referential indeterminacy problem, as well as the effect of the robot's autonomy on the acceptability of the system by naive users. Third, although *DAC-h3* can solve parts of the *Symbol Grounding Problem* (SGP), it still presupposes a symbolic concept of an entity which is given *a priori* to the system. Therefore, our contribution is more about the ability to proactively acquire multimodal information about these entities and linking them to linguistic labels that can be reused to express complex goal-oriented behavior later on.

We are currently extending the proposed architecture in the following ways. First, we are better integrating some of the available abilities within the interaction scenario as mentioned above. This will allow starting the knowledge acquisition process from scratch in a fully autonomous way. Second, we are considering to use more biologically plausible and/or computationally scalable models for some of the existing modules, namely the *action planning* and *action selection* modules. These are currently algorithmic implementations using predefined action plans. We want to replace it with an existing model of rule learning grounded in the neurobiology of the prefrontal cortex which can learn optimal action policies from experience to maximize long-term reward [105]. An interesting feature of this model for solving the SGP is that it relies on neural memory-units encoding sensorimotor contingencies with causal relationships learned through adaptive connections between them. An alternative solution is to use state-of-the-art AI planners such as the Planning Domain Definition Language (PDDL), where multi-agent planning extensions are particularly relevant in the context of social robotics [106]. Third, we are also integrating more low-level reactive control abilities through an acquired notion of a peri-personal space [107], where the robot will be able to optimize its own action primitives to maintain safety distances with aversive objects (e.g. a spider) in real time while executing reaching actions toward other objects. Finally, we are working on a self-exploration process to autonomously discover the area which is reachable by the robot, similarly to Jamone et al. [108], and subsequently employing this self-model and applying it to the human partner to estimate his/her reachability.

REFERENCES

- [1] S. Harnad, "The symbol grounding problem," *Physica D: Nonlinear Phenomena*, vol. 42, no. 1, pp. 335–346, 1990.
- [2] S. Coradeschi and A. Saffiotti, "An introduction to the anchoring problem," *Robotics and Autonomous Systems*, vol. 43, no. 2–3, pp. 85–96, 2003.
- [3] T. Taniguchi, T. Nagai, T. Nakamura, N. Iwahashi, T. Ogata, and H. Asoh, "Symbol emergence in robotics: a survey," *Advanced Robotics*, vol. 30, no. 11–12, pp. 706–728, 2016.
- [4] G. Metta et al., "The iCub humanoid robot: An open-systems platform for research in cognitive development," *Neural Networks*, vol. 23, no. 8, pp. 1125–1134, 2010.
- [5] P. Vogt, "The physical symbol grounding problem," *Cognitive Systems Research*, vol. 3, no. 3, pp. 429–457, 2002.
- [6] P. Vogt and F. Divina, "Social symbol grounding and language evolution," *Interaction Studies*, vol. 8, no. 1, pp. 31–52, 2007.
- [7] A. Cangelosi, "The grounding and sharing of symbols," *Pragmatics & Cognition*, vol. 14, no. 2, pp. 275–285, 2006.
- [8] S. Lallée et al., "Towards the synthetic self: making others perceive me as an other," *Paladyn Journal of Behavioral Robotics*, vol. 6, no. 1, 2015.
- [9] V. Vouloutsi et al., "Towards a Synthetic Tutor Assistant: The EASEL Project and its Architecture," in *International Conference on Living Machines*, 2016, pp. 353–364.
- [10] S. C. Levinson, "On the human interaction engine," in *Wenner-Gren Foundation for Anthropological Research, Symposium 134*. Berg, 2006, pp. 39–69.
- [11] U. Liszkowski, M. Carpenter, A. Henning, T. Striano, and M. Tomasello, "Twelve-month-olds point to share attention and interest," *Developmental science*, vol. 7, no. 3, pp. 297–307, 2004.
- [12] D. E. Berlyne, "A theory of human curiosity," *British Journal of Psychology. General Section*, vol. 45, no. 3, pp. 180–191, 1954.
- [13] M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll, "Understanding and sharing intentions: The origins of cultural cognition," *Behavioral and brain sciences*, vol. 28, no. 5, pp. 675–691, 2005.
- [14] D. Vernon, C. von Hofsten, and L. Fadiga, "Desiderata for developmental cognitive architectures," *Biologically Inspired Cognitive Architectures*, vol. 18, pp. 116–127, 2016.
- [15] J. Piaget, *The construction of reality in the child*. New York, NY, US: Basic Books, 1954.
- [16] L. Vygotsky, *Mind in Society: Development of Higher Psychological Processes*. Cambridge: Harvard University Press, 1978.
- [17] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, no. 3, pp. 143–166, 2003.
- [18] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun, "Towards robotic assistants in nursing homes: Challenges and results," *Robotics and Autonomous Systems*, vol. 42, no. 3, pp. 271–281, 2003.
- [19] I. R. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto, "An affective mobile robot educator with a full-time job," *Artificial Intelligence*, vol. 114, no. 1, pp. 95–124, 1999.
- [20] A. G. Di Nuovo, D. Marocco, S. Di Nuovo, and A. Cangelosi, "Autonomous learning in humanoid robotics through mental imagery," *Neural Networks*, vol. 41, pp. 147–155, 2013.
- [21] B. Adams, C. Breazeal, R. A. Brooks, and B. Scassellati, "Humanoid robots: a new kind of tool," *IEEE Intelligent Systems and their Applications*, vol. 15, no. 4, pp. 25–31, 2000.
- [22] K. Dautenhahn and A. Billard, "Bringing up robots or - the psychology of socially intelligent robots: From theory to implementation," in *Conference on Autonomous Agents*, 1999, pp. 366–367.
- [23] Y. Demiris, L. Aziz-Zadeh, and J. Bonaiuto, "Information Processing in the Mirror Neuron System in Primates and Machines," *Neuroinformatics*, vol. 12, no. 1, pp. 63–91, 2014.
- [24] A. Newell, *Unified theories of cognition*. Harvard University Press, 1990.
- [25] A. Newell, J. C. Shaw, and H. A. Simon, "Report on a general problem-solving program," *IFIP Congress*, pp. 256–264, 1959.
- [26] J. E. Laird, A. Newell, and P. S. Rosenbloom, "SOAR: An architecture for general intelligence," *Artificial Intelligence*, vol. 33, no. 1, pp. 1–64, 1987.
- [27] J. R. Anderson, *The Architecture of Cognition*. Harvard University Press, 1983.
- [28] R. A. Brooks, "Intelligence without representation," *Artificial Intelligence*, vol. 47, no. 1–3, pp. 139–159, 1991.
- [29] R. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal on Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986.
- [30] Á. Miklósi and M. Gácsi, "On the utilization of social animals as a model for social robotics," *Frontiers in Psychology*, vol. 3, no. 75, pp. 1–10, 2012.
- [31] K. Dautenhahn, "Socially intelligent robots: dimensions of human-robot interaction," *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, vol. 362, no. 1480, pp. 679–704, 2007.
- [32] P. Carruthers and P. Smith, *Theories of Theories of Mind*. Cambridge University Press, 1996.
- [33] E. Di Paolo and H. De Jaeger, "The interactive brain hypothesis," *Frontiers in Human Neuroscience*, vol. 6, no. 163, pp. 1–16, 2012.
- [34] T. Taniguchi, "Symbol emergence in robotics for long-term human-robot collaboration," in *IFAC Symposium on Analysis, Design, and Evaluation of Human-Machine Systems*, 2016, pp. 144–149.
- [35] B. Scassellati, "Investigating models of social development using a humanoid robot," in *IEEE International Joint Conference on Neural Networks*, 2003, pp. 2704–2709.
- [36] A. Stoytchev and R. Arkin, "Combining deliberation, reactivity, and motivation in the context of a behavior-based robot architecture,"

- in *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 2001, pp. 290–295.
- [37] M. Malfaz, Á. Castro-Gonzalez, R. Barber, and M. A. Salichs, “A Biologically Inspired Architecture for an Autonomous and Social Robot,” *IEEE Transactions on Autonomous Mental Development*, vol. 3, no. 3, pp. 232–246, 2011.
 - [38] M. Scheutz, G. Briggs, R. Cantrell, E. Krause, T. Williams, and R. Veale, “Novel mechanisms for natural human-robot interactions in the diarc architecture,” in *Proceedings of AAAI Workshop on Intelligent Robotic Systems*, 2013.
 - [39] P. F. M. J. Verschure, T. Voegtlin, and R. J. Douglas, “Environmentally mediated synergy between perception and behaviour in mobile robots,” *Nature*, vol. 425, no. 6958, pp. 620–624, 2003.
 - [40] P. F. M. J. Verschure, C. M. A. Pennartz, and G. Pezzulo, “The why, what, where, when and how of goal-directed choice: neuronal and computational principles,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 369, no. 1655, pp. 1–14, 2014.
 - [41] J.-Y. Puigbò, C. Moulin-Frier, and P. F. Verschure, “Towards self-controlled robots through distributed adaptive control,” in *Biomimetic and Biohybrid Systems*, 2016, pp. 490–497.
 - [42] P. Gärdenfors, *Conceptual Spaces. The Geometry of Thought*. Cambridge, MIT Press, 2000.
 - [43] A. Lieto, A. Chella, and M. Frixione, “Conceptual spaces for cognitive architectures: A lingua franca for different levels of representation,” *Biologically Inspired Cognitive Architectures*, vol. 19, pp. 1–9, 2017.
 - [44] C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen, “A large-scale model of the functioning brain,” *Science*, vol. 338, no. 6111, pp. 1202–1205, 2012.
 - [45] P. Blouw, E. Solodkin, P. Thagard, and C. Eliasmith, “Concepts as Semantic Pointers: A Framework and Computational Model,” *Cognitive Science*, vol. 40, no. 5, pp. 1128–1162, 2016.
 - [46] N. Krüger *et al.*, “Object–Action Complexes: Grounded abstractions of sensory–motor processes,” *Robotics and Autonomous Systems*, vol. 59, no. 10, pp. 740–757, 2011.
 - [47] L. Jamone, E. Ugur, A. Cangelosi, L. Fadiga, A. Bernardino, J. Piater, and J. Santos-Victor, “Affordances in psychology, neuroscience and robotics: a survey,” *IEEE Transactions on Cognitive and Developmental Systems*, 2016.
 - [48] D. Vernon, M. Beetz, and G. Sandini, “Prospection in Cognition: The Case for Joint Episodic-Procedural Memory in Cognitive Robotics,” *Frontiers in Robotics and AI*, vol. 2, no. 19, pp. 1–14, 2015.
 - [49] L. Steels, “The synthetic modeling of language origins,” *Evolution of Communication*, vol. 1, no. 1, pp. 1–34, 1997.
 - [50] F. Kaplan, “Semiotic schemata: Selection units for linguistic cultural evolution,” in *International Conference on Artificial Life*, 2000, pp. 372–381.
 - [51] C. Moulin-Frier, J. Diard, J.-L. Schwartz, and P. Bessière, “COSMO (“Communicating about Objects using Sensory-Motor Operations”): a Bayesian modeling framework for studying speech communication and the emergence of phonological systems,” *Journal of Phonetics*, vol. 53, pp. 5–41, 2015.
 - [52] S. Lallée *et al.*, “Cooperative human robot interaction systems: IV. Communication of shared plans with Naïve humans using gaze and speech,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 129–136.
 - [53] S. Lallée *et al.*, “Towards a platform-independent cooperative human robot interaction system: III an architecture for learning and executing actions and shared plans,” *IEEE Transactions on Autonomous Mental Development*, vol. 4, no. 3, pp. 239–253, 2012.
 - [54] M. Petit *et al.*, “The coordinating role of language in real-time multimodal learning of cooperative tasks,” *IEEE Transactions on Autonomous Mental Development*, vol. 5, no. 1, pp. 3–17, 2013.
 - [55] S. Calinon, F. D’halluin, E. L. Sauser, D. G. Caldwell, and A. G. Billard, “Learning and reproduction of gestures by imitation,” *IEEE Robotics & Automation Magazine*, vol. 17, no. 2, pp. 44–54, 2010.
 - [56] Y. Demiris and A. Meltzoff, “The robot in the crib: A developmental analysis of imitation skills in infants and robots,” *Infant and Child Development*, vol. 17, no. 1, pp. 43–53, 2008.
 - [57] S. Schaal, “Is imitation learning the route to humanoid robots?” *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
 - [58] K. Lee, Y. Su, T.-K. Kim, and Y. Demiris, “A syntactic approach to robot imitation learning using probabilistic activity grammars,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1323–1334, 2013.
 - [59] M. Ewerton, G. Maeda, J. Peters, and G. Neumann, “Learning motor skills from partially observed movements executed at different speeds,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 456–463.
 - [60] A. Billard and K. Dautenhahn, “Grounding communication in autonomous robots: an experimental study,” *Robotics and Autonomous Systems*, vol. 24, no. 1, pp. 71–79, 1998.
 - [61] A. Cangelosi, E. Hourdakis, and V. Tikhonoff, “Language acquisition and symbol grounding transfer with neural networks and cognitive robots,” in *IEEE International Joint Conference on Neural Network Proceedings*, 2006, pp. 1576–1582.
 - [62] D. Marocco, A. Cangelosi, K. Fischer, and T. Belpaeme, “Grounding action words in the sensorimotor interaction with the world: experiments with a simulated iCub humanoid robot,” *Frontiers in Neurorobotics*, vol. 4, no. 7, pp. 1–15, 2010.
 - [63] M. Quine, *Word and object*. MIT Press, 1960.
 - [64] P.-Y. Oudeyer, F. Kaplan, and V. V. Hafner, “Intrinsic motivation systems for autonomous mental development,” *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 2, pp. 265–286, 2007.
 - [65] C. Breazeal and B. Scassellati, “Infant-like social interactions between a robot and a human caregiver,” *Adaptive Behavior*, vol. 8, no. 1, pp. 49–74, 2000.
 - [66] S. Ivaldi *et al.*, “Object learning through active exploration,” *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 1, pp. 56–72, 2013.
 - [67] C. Moulin-Frier, S. M. Nguyen, and P.-Y. Oudeyer, “Self-Organization of Early Vocal Development in Infants and Machines: The Role of Intrinsic Motivation,” *Frontiers in Psychology*, vol. 4, no. 1006, 2014.
 - [68] O. C. Schrempf, U. D. Hanebeck, A. J. Schmid, and H. Worn, “A novel approach to proactive human-robot cooperation,” in *IEEE International Workshop on Robot and Human Interactive Communication*, 2005, pp. 555–560.
 - [69] I. Lutkebohle *et al.*, “The curious robot - structuring interactive robot learning,” in *IEEE International Conference on Robotics and Automation*, 2009, pp. 4156–4162.
 - [70] F. Broz *et al.*, “The ITALK Project: A Developmental Robotics Approach to the Study of Individual, Social, and Linguistic Learning,” *Topics in Cognitive Science*, vol. 6, no. 3, pp. 534–544, 2014.
 - [71] A. Cangelosi *et al.*, “Integration of action and language knowledge: A roadmap for developmental robotics,” *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 167–195, 2010.
 - [72] A. Antunes, L. Jamone, G. Saponaro, A. Bernardino, and R. Ventura, “From human instructions to robot actions: Formulation of goals, affordances and probabilistic planning,” in *IEEE International Conference on Robotics and Automation*, 2016, pp. 5449–5454.
 - [73] E. A. Krause, M. Zillich, T. E. Williams, and M. Scheutz, “Learning to Recognize Novel Objects in One Shot through Human-Robot Interactions in Natural Language Dialogues,” in *AAAI Conference on Artificial Intelligence*, 2014, pp. 2796–2802.
 - [74] J. Dias, W. C. Ho, T. Vogt, N. Beeckman, A. Paiva, and E. André, “I know what I did last summer: Autobiographic memory in synthetic characters,” in *International Conference on Affective Computing and Intelligent Interaction*, 2007, pp. 606–617.
 - [75] D. S. Syrdal, K. Dautenhahn, K. L. Koay, and W. C. Ho, “Views from within a narrative: Evaluating long-term human–robot interaction in a naturalistic environment using open-ended scenarios,” *Cognitive Computation*, vol. 6, no. 4, pp. 741–759, 2014.
 - [76] G. Sieber and B. Krenn, “Towards an episodic memory for companion dialogue,” in *International Conference on Intelligent Virtual Agents*, 2010, pp. 322–328.
 - [77] P. Fitzpatrick, G. Metta, and L. Natale, “Towards long-lived robot genes,” *Robotics and Autonomous Systems*, vol. 56, no. 1, pp. 29–45, 2008.
 - [78] S. Fanello *et al.*, “3D Stereo Estimation and Fully Automated Learning of Eye-Hand Coordination in Humanoid Robots,” in *IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 1028–1035.
 - [79] Z. Mathews, S. B. i Badia, and P. F. Verschure, “PASAR: An integrated model of prediction, anticipation, sensation, attention and response for artificial sensorimotor systems,” *Information Sciences*, vol. 186, no. 1, pp. 1–19, 2012.
 - [80] M. Sanchez-Fibla *et al.*, “Allostatic control for robot behavior regulation: a comparative rodent-robot study,” *Advances in Complex Systems*, vol. 13, no. 3, pp. 377–403, 2010.
 - [81] M. S. Fibla, U. Bernardet, and P. F. Verschure, “Allostatic control for robot behaviour regulation: An extension to path planning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1935–1942.
 - [82] S. Lallée and P. F. Verschure, “How? Why? What? Where? When? Who? Grounding Ontology in the Actions of a Situated Social Agent,” *Robotics*, vol. 4, no. 2, pp. 169–193, 2015.

- [83] V. Vouloutsi, K. Grechuta, S. Lallée, and P. F. Verschure, "The influence of behavioral complexity on robot perception," in *Conference on Biomimetic and Biohybrid Systems*, 2014, pp. 332–343.
- [84] U. Pattacini, F. Nori, L. Natale, and G. Metta, "An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1668–1674.
- [85] A. Parmiggiani *et al.*, "The Design of the iCub Humanoid Robot," *International Journal of Humanoid Robotics*, vol. 9, no. 4, 2012.
- [86] G. Pasquale, C. Ciliberto, L. Rosasco, and L. Natale, "Object identification from few examples by improving the invariance of a deep convolutional neural network," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 4904–4911.
- [87] X. Hinaut and P. F. Dominey, "Real-Time Parallel Processing of Grammatical Structure in the Fronto-Striatal System: A Recurrent Network Simulation Study Using Reservoir Computing," *PloS one*, vol. 8, no. 2, pp. 1–18, 2013.
- [88] X. Hinaut, M. Petit, G. Pointeau, and P. F. Dominey, "Exploring the acquisition and production of grammatical constructions through human-robot interaction with echo state networks," *Frontiers in Neurorobotics*, vol. 8, no. 16, pp. 1–17, 2015.
- [89] A.-L. Mealiar, G. Pointeau, P. Gardenfors, and P.-F. Dominey, "Construals of meaning: The role of attention in robotic language production," *Interaction Studies*, vol. 17, no. 1, pp. 41–69, 2016.
- [90] A. Damianou, C. H. Ek, L. Boorman, N. D. Lawrence, and T. J. Prescott, "A top-down approach for a synthetic autobiographical memory system," in *Biomimetic and Biohybrid Systems*, 2015, pp. 280–292.
- [91] A. C. Damianou and N. D. Lawrence, "Deep gaussian processes," in *Artificial Intelligence and Statistics Conference*, 2013, pp. 207–215.
- [92] D. Camilleri, A. Damianou, H. Jackson, N. Lawrence, and T. Prescott, "iCub Visual Memory Inspector: Visualising the iCub's Thoughts," in *Conference on Biomimetic and Biohybrid Systems*, 2016, pp. 48–57.
- [93] M. Zambelli, T. Fischer, M. Petit, H. J. Chang, A. Cully, and Y. Demiris, "Towards anchoring self-learned representations to those of other agents," in *Workshop on Bio-inspired Social robot Learning in Home Scenarios at IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.
- [94] M. Zambelli and Y. Demiris, "Online Multimodal Ensemble Learning using Self-learned Sensorimotor Representations," in *IEEE Transactions on Cognitive and Developmental Systems*, vol. 9, no. 2, 2017, pp. 113–126.
- [95] H. J. Chang and Y. Demiris, "Highly Articulated Kinematic Structure Estimation combining Motion and Skeleton Information," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017, to be published.
- [96] H. J. Chang, T. Fischer, M. Petit, M. Zambelli, and Y. Demiris, "Kinematic structure correspondences via hypergraph matching," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4216–4225.
- [97] T. Fischer and Y. Demiris, "Markerless Perspective Taking for Humanoid Robots in Unconstrained Environments," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 3309–3316.
- [98] M. Johnson and Y. Demiris, "Perceptual Perspective Taking and Action Recognition," *International Journal of Advanced Robotic Systems*, vol. 2, no. 4, pp. 301–308, 2005.
- [99] T. J. Prescott, N. Lepora, and P. F. M. J. Verschure, "A future of living machines? International trends and prospects in biomimetic and biohybrid systems," in *Conference on Bioinspiration, Biometrics and Bioreplication*, 2014.
- [100] P. Verschure, "Formal minds and biological brains II: From the mirage of intelligence to a science and engineering of consciousness," in *IEEE Intelligent Systems Trends and Controversies*, 2013, pp. 33–36.
- [101] G. Pointeau, M. Petit, and P. F. Dominey, "Successive Developmental Levels of Autobiographical Memory for Learning Through Social Interaction," *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 3, pp. 200–212, Sep. 2014.
- [102] M. Petit, T. Fischer, and Y. Demiris, "Lifelong Augmentation of Multi-Modal Streaming Autobiographical Memories," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 8, no. 3, pp. 201–213, 2016.
- [103] M. Petit, T. Fischer, and Y. Demiris, "Towards the emergence of procedural memories from lifelong multi-modal streaming memories for cognitive robots," in *Workshop on Machine Learning Methods for High-Level Cognitive Capabilities in Robotics at IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.
- [104] J.-Y. Puigbo, C. Moulin-Frier, V. Vouloutsi, M. Sanchez-Fibla, I. Hereros, and P. F. J. Verschure, "Skill refinement through cerebellar learning and human haptic feedback: An iCub learning to paint experiment," in *IEEE-RAS International Conference on Humanoid Robots*, 2015, pp. 447–452.
- [105] A. Duff, M. Sanchez Fibla, and P. F. M. J. Verschure, "A biologically based model for the integration of sensory-motor contingencies in rules and plans: A prefrontal cortex based extension of the Distributed Adaptive Control architecture," *Brain Research Bulletin*, vol. 85, no. 5, pp. 289–304, 2011.
- [106] D. L. Kovacs, "A Multi-Agent Extension of PDDL3.1," in *Workshop on International Planning Competition at the International Conference on Automated Planning and Scheduling*, 2012, pp. 19–27.
- [107] A. Roncone, M. Hoffmann, U. Pattacini, L. Fadiga, and G. Metta, "Peripersonal space and margin of safety around the body: learning tactile-visual associations in a humanoid robot with artificial skin," *PLoS ONE*, vol. 11, no. 10, pp. 1–32, 2016.
- [108] L. Jamone, M. Brandao, L. Natale, K. Hashimoto, G. Sandini, and A. Takanishi, "Autonomous online generation of a motor representation of the workspace for intelligent whole-body reaching," *Robotics and Autonomous Systems*, vol. 62, no. 4, pp. 556–567, 2014.