



This is a repository copy of *Towards a synthetic tutor assistant: The EASEL project and its architecture*.

White Rose Research Online URL for this paper:  
<http://eprints.whiterose.ac.uk/108443/>

Version: Accepted Version

---

**Proceedings Paper:**

Vouloutsis, V., Blancas, M., Zucca, R. et al. (17 more authors) (2016) Towards a synthetic tutor assistant: The EASEL project and its architecture. In: Lecture Notes in Computer Science. 5th International Conference, Living Machines 2016, 19-22 July 2016, Edinburgh, UK. Biomimetic and Biohybrid Systems, 9793 . Springer Verlag , pp. 353-364. ISBN 9783319424163

[https://doi.org/10.1007/978-3-319-42417-0\\_32](https://doi.org/10.1007/978-3-319-42417-0_32)

---

**Reuse**

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Towards a Synthetic Tutor Assistant: the EASEL project and its architecture

Vasiliki Vouloutsi<sup>1</sup>, Maria Blancas<sup>1</sup>, Riccardo Zucca<sup>1</sup>, Pedro Omedas<sup>1</sup>, Dennis Reidsma<sup>3</sup>, Daniel Davison<sup>3</sup>, Vicky Charisi<sup>3</sup>, Frances Wijnen<sup>3</sup>, Jan van der Meij<sup>3</sup>, Vanessa Evers<sup>3</sup>, David Cameron<sup>4</sup>, Samuel Fernando<sup>4</sup>, Roger Moore<sup>4</sup>, Tony Prescott<sup>4</sup>, Daniele Mazzei<sup>5</sup>, Michael Pieroni<sup>5</sup>, Lorenzo Cominelli<sup>5</sup>, Roberto Garofalo<sup>5</sup>, Danilo De Rossi<sup>5</sup>, and Paul F.M.J. Verschure<sup>1,2</sup>

<sup>1</sup> SPECS, N-RAS, DTIC, Universitat Pompeu Fabra (UPF), Barcelona, ES

<sup>2</sup> Catalan Institute of Advanced Studies (ICREA), Barcelona, ES

<sup>3</sup> Human Media Interaction / ELAN, University of Twente, NL

<sup>4</sup> Sheffield Robotics, University of Sheffield, UK

<sup>5</sup> University of Pisa, IT

{vicky.vouloutsi,maria.blancas,riccardo.zucca,pedro.omedas,paul.verschure}@upf.edu,  
{d.reidsma,v.charisi,d.p.davison,f.m.wijnen,j.vandermeij,v.evers}@utwente.nl,  
{d.s.cameron,s.fernando,r.k.moore,t.j.prescott}@sheffield.ac.uk,  
lorenzo.cominelli@for.unipi.it,{daniele.mazzei,michael.pieroni,roberto.garofalo,danilo.derossi}@centropiaggio.unipi.it

**Abstract.** Robots are gradually but steadily being introduced in our daily lives. A paramount application is that of education, where robots can assume the role of a tutor, a peer or simply a tool to help learners in a specific knowledge domain. Such endeavor posits specific challenges: affective social behavior, proper modelling of the learner’s progress, discrimination of the learner’s utterances, expressions and mental states, which, in turn, require an integrated architecture combining perception, cognition and action. In this paper we present an attempt to improve the current state of robots in the educational domain by introducing the EASEL EU project. Specifically, we introduce the EASEL’s unified robot architecture, an innovative Synthetic Tutor Assistant (STA) whose goal is to interactively guide learners in a science-based learning paradigm, allowing us to achieve such rich multimodal interactions.

**Keywords:** education, robotic tutor assistant, pedagogical models, cognitive architecture, Distributed Adaptive Control

## 1 Introduction

Robot technology has become so advanced that automated systems are gradually taking on a greater role in our society. Robots are starting to make their mark in many domains ranging from health-care and medicine, to automotive technologies, to service robots and robot companions.

Here our interest is in educational robots and, specifically, how automated technologies can assist learners with their intellectual growth in the classroom.

Despite being still disputed, the use of robots in education has been shown to positively affect students' concentration, learning interest and academic achievements [1]. The introduction of assisting robots as part of a course has been proved effective for integration, real-world issues, interdisciplinary work as well as critical thinking [2]. Moreover, educational robots can be flexible, as they can assume the role of mere tools [3], peers [4, 5] or even tutors [6]. Although the preferred characterization is not yet conclusive [7], when used as a tutor or a peer, thus implying a continuous robot–learner interaction, the robot's design and behavior become a crucial aspect [6, 8, 9]. Typically, robots in education are employed in scenarios ranging from technical education [10] (usually related to robotics or technology), to science [11] and learning of a foreign language [5], just to cite a few – see also [12] for a recent review and discussion about the use of robots in education.

Although a comprehensive examination of pedagogical theories in educational robotics is still lacking, two main approaches have been mostly influential. On one side, Piaget's theory of *constructivism*, which defines knowledge as the process in which an individual creates a meaning out of his own experiences [13]. On the other side, Papert's theory of *constructionism*, stating that learning is the result of building knowledge structures through progressive internalization of actions and conscious engagement through making [14]. More recently, the work of the Russian psychologist Lev Vygotsky gained more and more attention, as it introduced the principle of *scaffolding* and the one of *Zone of Proximal Development (ZPD)* [15]. The former refers to the usage of tools or strategies providing help, whereas the latter corresponds to the distance between what a learner can do by himself and what he may do under the guidance of an effective mediator. All these pedagogical approaches are highly relevant to robotic applications in education, and a review of related studies can be found in [16].

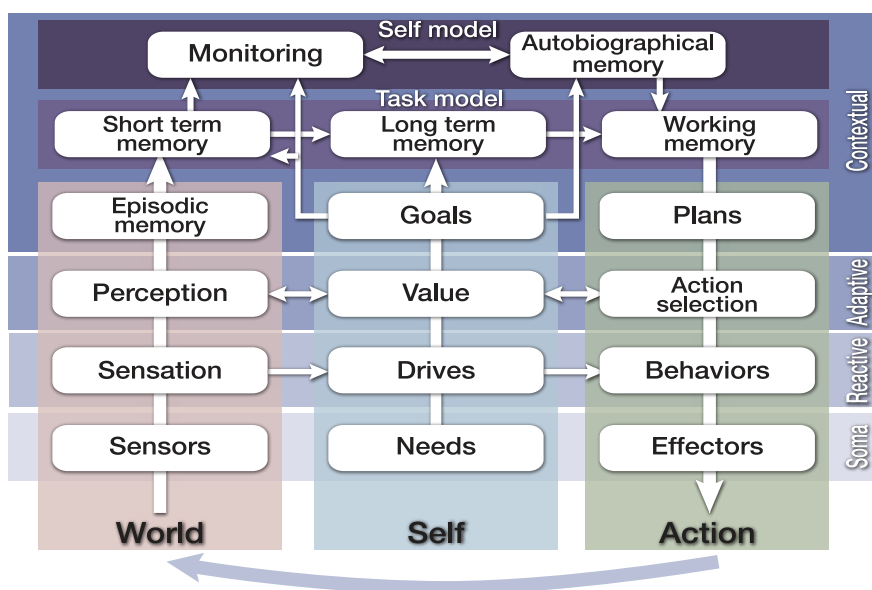
The work we present here constitutes an effort to move one step forward in the domain of robots in education by introducing the EASEL EU–project. EASEL is the acronym of *Expressive Agents for Symbiotic Education and Learning* and it is a collaborative project aimed to explore and develop a theoretical understanding of the Human-Robot Symbiotic Interaction (HRSI) realized in the domain of tutoring. The final outcome of EASEL is the delivery of an innovative Synthetic Tutor Assistant (STA), whose goal is to interactively guide learners (e.g., children in the age 8–11) using a science-based learning paradigm. The main focus of this paper is in describing the underlying STA's architecture that will allow the agent to interact with a human user in an educational task, whereas the theoretical approach and expected impact can be found in an accompanying paper (Reidsma et al. [in preparation]).

The paper's structure is as follows: in Section 2 we present the Distributed Adaptive Control (DAC) model of human and animal learning that serves as both a pedagogical model and the control architecture of the STA. In Section 3 we present the developed educational scenarios that serve as both the test case

of the architecture as well as an evaluation method. Finally, in Section 4 we introduce the overall architecture and present its individual components.

## 2 The DAC architecture and pedagogical model

The STA's main goal is to guide the learner through the science-based educational scenario in order to maximize learning. Based on the perception of the social, communicative and educational context, the robot will respond accordingly within the educational scenario and the specific learning task. The STA's reasoning and memory components need to continuously extract relevant knowledge from sequences of behavioural interactions over prolonged periods to learn a model of the user and adapt its actions to the partner. To successfully interact with the user, the STA thus requires an integrated architecture combining perception, cognition, and action. Here, we adopt the Distributed Adaptive Control (DAC) cognitive architecture (Figure 1) as the basis to control the STA's behavior.



**Fig. 1.** Schematic illustration of the Distributed Adaptive Control architecture (DAC) and its four layers: somatic, reactive, adaptive and contextual. Across the layers we can distinguish three main functional columns of organization: world related (exosensing, red), self related (endosensing, blue) and the interface to the world through action (green). The arrows indicate the flow of information. Adapted from [17].

DAC [17–19] is a robot-based neuronal model of perception, cognition and behavior that is a standard in the domains of new artificial intelligence and behavior-based robotics. It is biologically constrained and fully grounded since it autonomously generates representations of its primary sensory input [18, 20]. The DAC architecture is organized around four tightly coupled layers: Soma, Reactive, Adaptive and Contextual. Across these layers, three functional columns of organization can be distinguished: exosensing, defined as the sensation and perception of the world, endosensing which detects and signals the states of the self, and the interface to the world through action.

The Soma represents the body itself and the information acquired from sensation, needs and actuation. In the Reactive Layer, predefined sensorimotor reflexes are triggered by low complexity signals and are coupled to specific affective states of the agent. The Adaptive Layer extends these sensorimotor loops with acquired sensor and action states, allowing the agent to escape the predefined reflexes through learning. Finally, the Contextual Layer develops the state-space acquired by the Adaptive Layer to generate behavioral plans and policies that can be expressed through actions.

Furthermore, DAC postulates that learning is organized along a similar hierarchy of complexity. In order to learn and consolidate new knowledge, the learner undergoes three different learning stages: *resistance*, *confusion* and *abduction*. Resistance is the mechanism that results from defending one’s own world-model and is highly related to maintaining the balance of one’s feeling of agency. All agents possess an internalized world-model, which they need to reconsider when exposed to new knowledge or experience. However, learners tend to hold overly optimistic and confused views about their level of knowledge: those with good expertise have a tendency to underestimate their overall capabilities, whereas those who don’t, tend to overestimate their abilities [21]. Resistance is what consequently leads to the state of confusion, which actually generates the necessity to resolve the problem and learn through re-adapting. Adjusting to individual’s skills and progress helps the process of learning acquisition: it is therefore essential to maintain a challenging enough task by adjusting the level of confusion to the skills and progress of the learner. Monitoring, controlling and adjusting confusion is what we define as *shaping the landscape of success*. Such approach is comparable to instructional scaffolding, a learning process intended to help the student to cross what Vygotsky called the Zone of Proximal Development [15]. Confusion needs to be controlled such that the task to learn is not too easy to become boring and, at the same time, it should not be too challenging, thus leading to a complete loss of motivation or development of learned helplessness [22]. The learner needs to believe that he can be effective in controlling the relevant events within the learning process [23]. Confusion is needed in order to discover and generate theories to assess them later, that is, to be able to perform abduction, which is the very process of acquiring the new knowledge.

In the EASEL project, the role of DAC is thus twofold: on the one hand, it acts as a control system to define and drive the STA’s actions; on the other hand, it serves as the pedagogical model through which learning can be achieved.

### 3 Science-based educational scenarios

For the EASEL project we designed two interaction scenarios in real-life settings based on inquiry-based learning tasks. Typically, inquiry-based learning tasks involve active exploration of the world, asking questions, making discoveries and testing hypotheses. The first scenario aims at teaching children about physics concepts based on the Piagetian balance-beam experiments. The second scenario is meant to help children learn about healthy living and physical exercise. These two situations allow to exploit two different perspectives: a formal teaching scenario compared to a more “voluntary free exploration”, and a language oriented versus more bodily involvement.


#### 3.1 The balance beam

The *balance beam problem* was first described by Inhelder and Piaget to characterize and explain children’s stages of cognitive development [24]. Following the Piagetian work, Siegler developed a methodology which allowed him to classify children’s cognitive developmental stages on the base of four rules of increasing complexity that children of different ages would apply while solving the balance beam task (Figure 2) [25, 26].

Briefly, in the balance beam scenario different numbers of weights are placed at varying distances from the fulcrum on the equally spaced pegs positioned on both arms of the scale. Children explore the physics of the balance problem using tangible materials and guided by an artificial agent (e.g., a robot or a virtual avatar) that serves as the bodily manifestation of the STA. Children are then asked to predict the behavior of the beam given the configuration provided: if it will stay in equilibrium, tip to the left or tip to the right. To succeed in this task children have to identify the relevant physical concepts (i.e., weight and distance) and understand the underlying multiplicative relation between the two variables (i.e., the “torque rule”). The goal of the interaction is that the child learns about balance and momentum by going through a series of puzzle tasks with the balance beam. The artificial agent is there to encourage the students, to help them get through the different tasks and to provide feedback; thus, learning improves by constantly monitoring the learner’s progresses.

#### 3.2 Healthy-Living

The second interaction scenario designed is that of healthy-living. This scenario involves a synthetic agent assisting learners in an inquiry-based learning task about the benefits of physical exercise. Specifically, the children will investigate the effects of different types of exercise on a number of physiological variables (e.g., heart rate, blood pressure, respiration rate, etc.). Two types of sessions are considered: in the first one, the artificial agent will encourage the children to perform exercises at varying speeds. It will then provide information about the outcome of the exercise in a friendly, accessible way, either via voice or through a video display allowing the children to have immediate feedback about their



**Fig. 2.** Schematic illustration of the four rules assessed by Siegler [25]. At each developmental stage one or both dimensions (i.e., weight and distance) are considered. For instance, Rule I exclusively considers the weight, whereas Rule IV considers both weights and distance from the fulcrum.

own state. In the second type of session, the child sits down with a prepared worksheet about healthy-living, which is used to prompt a spoken dialogue with the artificial agent. The child reads through the text on the worksheet, which comprises instructions and questions to the artificial agent about the previous interaction. For instance, it explains how the physiological values are related to other forms of energy, such as the calories contained in different foods, and how much energy is burnt during an exercise. The worksheet prompts the child to start off the interaction providing cues about the kind of things that can be said to the artificial agent.

## 4 Architecture Overview

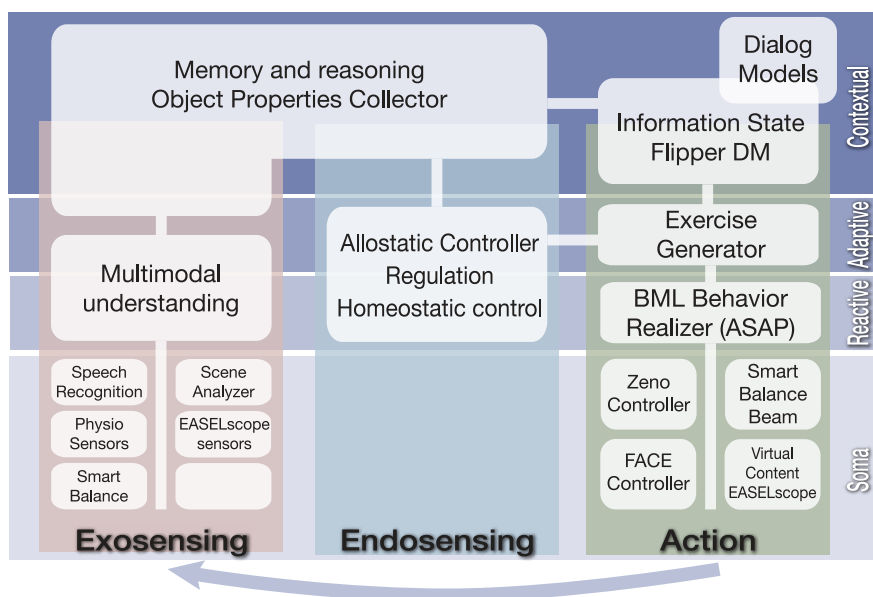
The proposed interaction scenarios consist of a humanoid robot, a handheld device (e.g., tablet) called the EASELscope and in the case of the balance beam task, the Smart Balance Beam (SBB), a motorized scale equipped with sensors to detect the object’s weight and position. Both scenarios imply a detailed interaction between the learner and the robot which create a series of challenges such as effective social behaviour to get a good social relation with the child, proper modelling of the students’ learning progress, various kinds of perception of the child’s utterances, expressions, bodily and mental states, as well as bodily and facial expressions from the robot.

Rich multimodal interactions require an integrated architecture which combines perception, reasoning, and action (Section 2). Such integration is more than just technically running modules side-by-side and exchanging messages through the YARP middleware communication platform [27]. It also requires alignment on the content, parameters, knowledge bases, and rule systems of all modules. In what follows, we present the unified EASEL architecture that makes all the interaction possible.

All modules communicate and exchange messages through YARP via named communication channels. This middleware platform allows us to not only distribute the system over different machines, but also permits abstraction from different Operating Systems.

#### 4.1 Modules and mapping to the DAC architecture

The integrated EASEL architecture is a practical incarnation of the conceptual architecture of the Distributed Adaptive Control (DAC) theory of the design principles underlying perception, cognition and action. Each implemented module in the EASEL architecture can be mapped to one or more of the core components of the DAC model that they embody and the specific modules are schematically illustrated in Figure 3.



**Fig. 3.** Overview of the EASEL architecture, where each implemented module is mapped to the core components of the DAC architecture.

The Speech recognition (ASR module), the SceneAnalyzer, the PhysioReader and the EASELScope sensors embed the exosensing component of the Soma layer through which the states of the world are acquired and the internal drives are established. More precisely, the ASR module is based on the open-source Kaldi speech recognition toolkit [28] with an EASEL specific vocabulary, language model and recognition grammar.

The SceneAnalyzer builds upon several other libraries to deliver integrated recognition of multimodal features of the users and their behaviour [29, 30]. The physiological signal acquisition module uses non-obtrusive and robust methods for obtaining information about the users physiological state: by integrating sensors in the robot or in the EASELScope tablet, information can be unobtrusively obtained without sensors worn or strapped to the body of user. The



EASELScope sensors allow detection of the current state of the balance beam allowing the EASEL system to respond to the actions of the user with respect to the learning materials.

In the reactive layer, ASAPRealizer module [31, 32] is responsible for the choreography of the behavior (verbal and non-verbal) of the STA using the generic robot-independent Behavior Markup Language (BML). We also use an easily configurable XML binding between the BML and the motion primitives of each robotic platform (that serves as the physical instantiation of the STA). Such approach abstracts away from specific motor control by exposing more general behaviour specifications to the dialog manager and provides generalization across embodiments. The ASAPRealizer maps to the behavior component of the DAC architecture by directly controlling the actuators of the somatic layer.

The Allostatic Control (AC) module currently implemented in the EASEL architecture embraces both the Reactive and the Adaptive layers of DAC. An homeostatic controller continuously classifies the current state of each drive by sending fast requests for corrective actions to keep drives within optimal boundaries. The allostatic controller maintains consistency between drives in an adaptive way by assigning priorities to the different drives and making the appropriate corrections to maintain coherence (e.g., by adapting the difficulty of the task to the learners behavior). The learning algorithms of the allostatic controller [33, 34] allow the STA to adapt its drives and homeostatic boundaries to a specific student’s behaviour and skills. Successful interactions (i.e., contextual cues and actions) are then stored as memory segments in the Object Properties Collector (OPC) module to build a model of the user.

The Exercise Generator contains the collection of learning exercises with all their different properties and difficulty levels. It selects the appropriate exercise given the current state of the tutoring model, the student model, and the output of the AC. This information is shared with the Flipper Dialog Manager to allow the robotic assistant to discuss with the child the progress of the exercise.

The Object Properties Collector (OPC) embodies the Contextual Layer’s memory components of DAC. At this stage of development, the OPC implements the memory for events that can be stored and distributed to the other STA’s modules as well as its short-term memory component. At each instant of the interaction, the OPC can temporarily retain ongoing perceptions, actions and values (i.e., outcomes of the current interaction) as segments of memory (relations). These relations allow the definition of rules for specific interactions that can be further stored as long-term memories if a high-level goal is successfully achieved. Between the OPC and the sensors lies the multimodal understanding module, a light and simple interpreter of speech, emotions and speaker’s probabilities, which simplifies the requirements for the Flipper Dialog Manager’s scripts. Flipper offers flexible dialog specification via information state and rule-based templates to trigger information state changes as well as behaviour requests [35].

The robots (Zeno (Hanson Robotics, Hong Kong) and FACE [36]), the virtual robot avatars and the EASEL-scope hidden state visualizer correspond to

the DAC effectors and represent the main interface of the STA with the world (Figure 4). The EASELscope offers an augmented reality (AR) interface that allows the learner to interact with the task materials. It can be used to present extra information to the child about the learning content. This allows the system to vary between different ways of scaffolding the learning of the user. For instance, using the EASELscope we can present “hidden information” about the balance beam, such as the weights of the pots, or the forces acting on the arms of the beam; both are types of scaffolds in learning that would not be possible without the scope (Figure 4).



**Fig. 4.** The Zeno robot, the EASELscope and the Smart Balance Beam (SBB). Zeno acts as the embodiment of the STA and continuously interacts with the learner. Information about the task is displayed as augmented reality on top of the physical SBB.

## 5 Discussion and conclusions

In this paper we presented the DAC based EASEL architecture, designed to guide learners through learning in two science-based educational scenarios. Each module within the framework has been integrated in a cohesive setup and the configuration options, models, behaviour repertoires and dialog scripts will allow us to validate the EASEL system through specific experiments with child-robot interaction in the proposed learning scenarios.

The way the architecture is organized gives us three key advantages: scalability, configurability and abstraction. This allows us to easily add sensory components with negligible changes to the main core of the system: it is sufficient to add the input to the multimodal understanding module that, in turn, will store the new information with an appropriate format in the OPC module. Furthermore, all modules are fully configurable: for instance, we can add new behaviors (ASAPRealizer), drives (Allostatic Control) as well as dialogues (Flipper) in an easy way with the usage of configuration files such as XML scripts. Thus, any

additional implementation for the needs of the EASEL architecture (in terms of scenarios or sensory inputs) can be done in a flexible way. Finally, the proposed architecture permits abstraction from the physical manifestation of the STA, in a way that using the same scenario, we can choose the robotic platform (or even avatar) with small changes in the main core of the system.

At this stage, we are now ready to start validating our educational architecture and focus on concrete long-term studies on human-robot symbiotic interactions in learning tasks. For instance, by looking at the types of hypotheses that the child frames and the outcome in solving exercises, the STA can continuously build and refine a model of the student’s understanding of the task. By varying the exercises, the scaffolding provided by the robot or the hidden-states information provided through the EASELScope, the STA can explore the most effective learning strategies for a specific task. By modifying the social and conversational strategies of the robot, the STA can extract the best “personality” and style of the robot leading to a good relationship between the child and the robot itself and their impact on the overall learning process.

## Acknowledgments

This work is supported by grants from the European Research Council under the European Union’s 7th Framework Programme FP7/2007-2013/ERC grant agreement n. 611971 (EASEL) and n. 341196 (CDAC) to Paul F. M. J. Verschure.

## References

1. J.-H. Han, M.-H. Jo, V. Jones, and J.-H. Jo, “Comparative study on the educational use of home robots for children,” *Journal of Information Processing Systems*, vol. 4, no. 4, pp. 159–168, 2008.
2. R. D. Beer, H. J. Chiel, and R. F. Drushel, “Using autonomous robotics to teach science and engineering,” *Communications of the ACM*, vol. 42, no. 6, pp. 85–92, 1999.
3. F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli, “The e-puck, a robot designed for education in engineering,” in *Proceedings of the 9th conference on autonomous robot systems and competitions*, vol. 1, pp. 59–65, IPCB: Instituto Politécnico de Castelo Branco, 2009.
4. F. Wijnen, V. Charisi, D. Davison, J. van der Meij, D. Reidsma, and V. Evers, “Inquiry learning with a social robot: can you explain that to me?,” in *Proceedings of New Friends 2015: the 1st international conference on social robotics in therapy and education* (M. Heerink and M. de Jong, eds.), (Almere, the Netherlands), pp. 24–25, Windesheim Flevoland, 2015.
5. T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro, “Interactive robots as social partners and peer tutors for children: A field trial,” *Human-computer interaction*, vol. 19, no. 1, pp. 61–84, 2004.
6. M. Saerbeck, T. Schut, C. Bartneck, and M. D. Janse, “Expressive robots in education: varying the degree of social supportive behavior of a robotic tutor,” in

- Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 1613–1622, ACM, 2010.
7. N. Shin and S. Kim, “Learning about, from, and with robots: Students’ perspectives,” in *Robot and Human interactive Communication, 2007. RO-MAN 2007. The 16th IEEE International Symposium on*, pp. 1040–1045, IEEE, 2007.
  8. V. Vouloutsi, M. B. Munoz, K. Grechuta, S. Lallee, A. Duff, J.-Y. P. Llobet, and P. F. M. J. Verschure, “A new biomimetic approach towards educational robotics: the distributed adaptive control of a synthetic tutor assistant,” *New Frontiers in Human-Robot Interaction*, p. 22, 2015.
  9. M. Blancas, V. Vouloutsi, K. Grechuta, and P. F. M. J. Verschure, “Effects of the robots role on human-robot interaction in an educational scenario,” in *Biomimetic and Biohybrid Systems*, pp. 391–402, Springer, 2015.
  10. T. Balch, J. Summet, D. Blank, D. Kumar, M. Guzdial, K. O’Hara, D. Walker, M. Sweat, G. Gupta, S. Tansley, J. Jackson, M. Gupta, M. Nur Muhammad, S. Prashad, N. Eilbert, and A. Gavin, “Designing personal robots for education: Hardware, software, and curriculum,” 2008.
  11. K. Highfield, J. Mulligan, and J. Hedberg, “Early mathematics learning through exploration with programmable toys,” in *Proceedings of the Joint Meeting of PME 32 and PME-NA*, pp. 169–176, Citeseer, 2008.
  12. O. Mubin, C. J. Stevens, S. Shahid, A. Al Mahmud, and J.-J. Dong, “A review of the applicability of robots in education,” *Journal of Technology in Education and Learning*, vol. 1, pp. 209–0015, 2013.
  13. J. Piaget and B. Inhelder, *The psychology of the child*. Basic Books, 1972.
  14. S. Papert and I. Harel, “Situating constructionism,” *Constructionism*, vol. 36, pp. 1–11, 1991.
  15. L. S. Vygotsky, *Mind in society: The development of higher psychological processes*. Harvard university press, 1980.
  16. V. Charisi, D. Davison, F. Wijnen, J. Van Der Meij, D. Reidsma, T. Prescott, W. Van Joolingen, and V. Evers, “Towards a child-robot symbiotic co-development: A theoretical approach,” in *AISB Convention 2015*, The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB), 2015.
  17. P. F. M. J. Verschure, “Distributed adaptive control: A theory of the mind, brain, body nexus,” *Biologically Inspired Cognitive Architectures*, 2012.
  18. 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. October, pp. 620–624, 2003.
  19. P. F. M. J. Verschure, C. M. Pennartz, and G. Pezzulo, “The why, what, where, when and how of goal-directed choice: neuronal and computational principles,” *Phil. Trans. R. Soc. B*, vol. 369, no. 1655, p. 20130483, 2014.
  20. G. Maffei, D. Santos-Pata, E. Marcos, M. Sánchez-Fibla, and P. F. M. J. Verschure, “An embodied biologically constrained model of foraging: from classical and operant conditioning to adaptive real-world behavior in dac-x,” *Neural Networks*, vol. 72, pp. 88–108, 2015.
  21. J. Kruger and D. Dunning, “Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments.,” *Journal of personality and social psychology*, vol. 77, no. 6, p. 1121, 1999.
  22. L. Y. Abramson, M. E. Seligman, and J. D. Teasdale, “Learned helplessness in humans: critique and reformulation.,” *Journal of abnormal psychology*, vol. 87, no. 1, p. 49, 1978.
  23. M. E. Seligman, “Learned helplessness,” *Annual review of medicine*, vol. 23, no. 1, pp. 407–412, 1972.

24. B. Inhelder and J. Piaget, *The growth of logical thinking from childhood to adolescence: An essay on the construction of formal operational structures*. Basic Books, 1958.
25. R. S. Siegler, “Three aspects of cognitive development,” *Cognitive psychology*, vol. 8, no. 4, pp. 481–520, 1976.
26. R. S. Siegler, S. Strauss, and I. Levin, “Developmental sequences within and between concepts,” *Monographs of the Society for Research in Child Development*, pp. 1–84, 1981.
27. G. Metta, P. Fitzpatrick, and L. Natale, “Yarp: yet another robot platform,” *International Journal on Advanced Robotics Systems*, vol. 3, no. 1, pp. 43–48, 2006.
28. D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, *et al.*, “The kaldi speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584, IEEE Signal Processing Society, 2011.
29. A. Zarakı, D. Mazzei, M. Giuliani, and D. De Rossi, “Designing and evaluating a social gaze-control system for a humanoid robot,” *Human-Machine Systems, IEEE Transactions on*, vol. 44, no. 2, pp. 157–168, 2014.
30. A. Zarakı, D. Mazzei, N. Lazzeri, M. Pieroni, and D. De Rossi, “Preliminary implementation of context-aware attention system for humanoid robots,” in *Biomimetic and Biohybrid Systems*, pp. 457–459, Springer, 2013.
31. D. Reidsma and H. van Welbergen, “Asaprealizer in practice—a modular and extensible architecture for a bml realizer,” *Entertainment computing*, vol. 4, no. 3, pp. 157–169, 2013.
32. H. Van Welbergen, R. Yaghoubzadeh, and S. Kopp, “Asaprealizer 2.0: the next steps in fluent behavior realization for ecas,” in *Intelligent Virtual Agents*, pp. 449–462, Springer, 2014.
33. V. Vouloutsi, S. Lallée, and P. F. M. J. Verschure, “Modulating behaviors using allostatic control,” in *Biomimetic and biohybrid systems*, pp. 287–298, Springer, 2013.
34. S. Lallée, V. Vouloutsi, S. Wierenga, U. Pattacini, and P. F. M. J. Verschure, “Efaa: a companion emerges from integrating a layered cognitive architecture,” in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pp. 105–105, ACM, 2014.
35. M. ter Maat and D. Heylen, “Flipper: An information state component for spoken dialogue systems,” in *Intelligent Virtual Agents*, pp. 470–472, Springer, 2011.
36. N. Lazzeri, D. Mazzei, A. Zarakı, and D. De Rossi, “Towards a believable social robot,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8064 LNAI, pp. 393–395, 2013. cited By 3.