This is a repository copy of *Cognitive architecture as a service: scaffolded integration of heterogeneous models through event streams*.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/195293/

Version: Accepted Version

Proceedings Paper:

Jimenez-Rodriguez, A., Robillard, J. and Prescott, A. orcid.org/0000-0003-4927-5390 (2022) Cognitive architecture as a service: scaffolded integration of heterogeneous models through event streams. In: Hunt, A., Vouloutsi, V., Moses, K., Quinn, R., Mura, A., Prescott, T. and Verschure, P.F.M.J., (eds.) Biomimetic and Biohybrid Systems : 11th International Conference, Living Machines 2022, Virtual Event, July 19–22, 2022, Proceedings. 11th International Conference, Living Machines 2022, 19-22 Jul 2022, Virtual Event. Lecture Notes in Computer Science, LNAI 13548 . Springer , pp. 345-353. ISBN 9783031204692

https://doi.org/10.1007/978-3-031-20470-8_34

This version of the contribution has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-031-20470-8_34. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

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 including the URL of the record and the reason for the withdrawal request.

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

Cognitive architecture as a service: scaffolded integration of heterogeneous models through event streams *

Alejandro Jimenez-Rodriguez[1,2][0000-0001-7172-1794], Julie Robillard[3][0000-0001-5765-4927], and Tony Prescott[1,2][0000-0003-4927-5390]

[1] University of Sheffield, Sheffield, UK
[2] Sheffield Robotics, Sheffield, UK
[3] The University of British Columbia, Vancouver BC, Canada
email{a.jimenez-rodriguez, t.j.prescott}@sheffield.ac.uk

Abstract. The development of cognitive architectures for biomimetic robots can benefit from the seamless integration of computational models that capture some of the brain's capacity to co-ordinate adaptive behavior. Such integration could take advantage of recent advances in distributed systems technology to support the communication between models, however, a communication protocol general enough to allow for heterogeneity, yet, simple enough to be practical and widely used, remains elusive. In this work we propose a solution based on a scaffolded structure that provides constraints for the different models to satisfy. Within this paradigm, the models do not interact among themselves but communicate using event sourcing technology supported by the open source stream processing platform *Apache Kafka*. This design allows the integration of brain-based models without having to specify module-to-module interfaces. At the same time, the robot acts as a consumer and producer of events through the *Neurorobotic Platform* (NRP) (part of the Human Brain Project's *EBrains* platform), meaning that the cognitive architecture has the potential to integrate components provided by a growing community of computational neuroscientists, and to be integrated with different robot platforms. In this paper we present this approach, which we term *Cognitive architecture as a Service* (CaaS), which is further motivated by the goal of creating assistive robots for human care settings. We also describe some early results, based on the *MiRo-e* robot platform, aimed at the development and evaluation of brain-based control for applications in this setting.

Keywords: Cognitive Architectures · Cognitive architectures as a Service · Distributed Systems · brain-based robots · socially-assistive robots.

1 Introduction

Socially assistive robots (SARs) are designed to support communication and interaction with people in human social and care settings, with the goal of sup-

---

* Supported by organization the Human Brain Project

porting engaging and valued short- and long-term human-robot interactions, and to act as catalysts for human-human interaction. The potential value of social robots in care settings has been demonstrated in a variety of studies, both in pediatric populations [10,14] and in older adults [16]. Recent work also indicates the potential of social robots that integrate some of the affective components of human social interaction, based on their capacity to provide emotional support that can lead to improved feelings of self-worth [10, 20]. The capacity of social robots to promote human-human interaction has also been demonstrated, for example, with the *Paro* animal-like robot which has been found to facilitate group interaction between adults with dementia [17, 23].

In order to safely interact and effectively communicate with people, SARs must embed a control system, or 'cognitive architecture' [13], that includes many human-like functional capacities including verbal and non-verbal communication, person and object detection and recognition, scene analysis and world knowledge, action and interaction planning. Whilst existing SARas differ substantially in the design and configuration of this architecture, many have been directly inspired by human and animal psychology, and several have integrated ethologically-inspired or computational neuroscience models [5–8, 15, 18, 19, 24]. This prior work demonstrates the potential for robots with biomimetic control systems to be useful in care settings. Indeed, the capacity of brain-inspired control to generate life-like behavior could be an advantage in providing interactions that people find rewarding and engaging.

In order to provide the infrastructure for a biomimetic robot cognitive architecture that includes heterogenous components we propose a *Cognitive architecture as a Service* (CaaS) paradigm which leverages modern distributed systems methods. Specifically, the CaaS is designed to serve as an scaffold for the integration of both engineered and brain-inspired components such that they can interact to produce continuous and appropriate behaviour, despite operating on different time-scales, and with radically different forms of internal communication and representation.

Thanks to this approach, the cognitive architecture is brought to the foreground of the communication infrastructure, providing an unified interface for models to communicate based upon functional considerations. The concept of a scaffold allow us to constrain successful existing distributed technologies with domain specific design choices [9, 27], therefore reducing the degrees of freedom available for the component models to communicate. The scaffold, in this case, is informed by the architecture of the mammalian brain [2].

The rest of the paper is structured as follows: in the next section we provide a conceptual description of the CaaS based upon the low level requirements of an autonomous SAR. We then provide a simple example that illustrates the main aspects of the architecture. Finally, we provide some conclusions and outlook.

2 The architecture

Previous approaches for the integration of heterogenous models into cognitive architectures have used bespoke distributed systems technologies for message passing [1, 4] between brain-inspired computational models. These particular architectural decisions constrain the design of the cognitive architecture itself. Particularly challenging is the fact that sometimes the input of one model comes from the output of another one, potentially developed in a different programming language, often using a different data representation.

In contrast, in the current work, the CaaS is designed using an event sourcing system based on open source *Apache Kafka* technology [12]. Kafka is an event stream platform that provides the necessary infrastructure for producing, consuming and storing events in distributed system. We also use the Neurorobotics platform (NRP) as part of the backbone [9]. The NRP consists of a set of *engines* that interact through a central core through the passing of messages or datapacks in order to execute simulations that control virtual or physical robots.

In using a set of event streams as the backbone of our architecture, we remove the need of the models to communicate with each other. Instead, the models produce and consume events in a set of predefined streams that are part of a pre-existing scaffold. This approach makes our system more flexible and extensible as it places fewer constraints on the design of interactions between modules. Figure 1 shows the design of the event stream of a first-generation CaaS inspired by the cognitive architecture of the mammalian brain. This CaaS builds on an earlier design developed for the MiRO-e animal-like robot an integrates a model of the vertebrate basal ganglia as a core action selection and conflict resolution mechanism [18]. The robot itself acts as a producer and consumer in the architecture.

The main conceptual aspects of this design are:

1. The raw perceptions enter the engine through datapacks that are collected and exposed as perception streams to be consumed and modified by the different models. Inspired by mammalian vision, perception splits into a "what" stream and a "where" stream where additional models can input object recognition, salience and location information. The streams are joined in the motivational controller which filters the perceptions according to the internal state of the robot [3]. This controller then selects between three phases of behaviour (seek, pursue, consume) according to the presence and distance to a goal. Seek will activate exploration sequences in the basal ganglia, pursue will activate approach or avoid actions and consume will execute predefined sequences like grasping.
2. The basal ganglia performs the actions selection based on the different motivational streams that are active in the robot, generally favouring one stream over the others, and produces to a stream of actions which represents the current action.
3. Part of the philosophy of Kafka is that the history of events that happened to the robot or the actions that it has performed are available (up to some

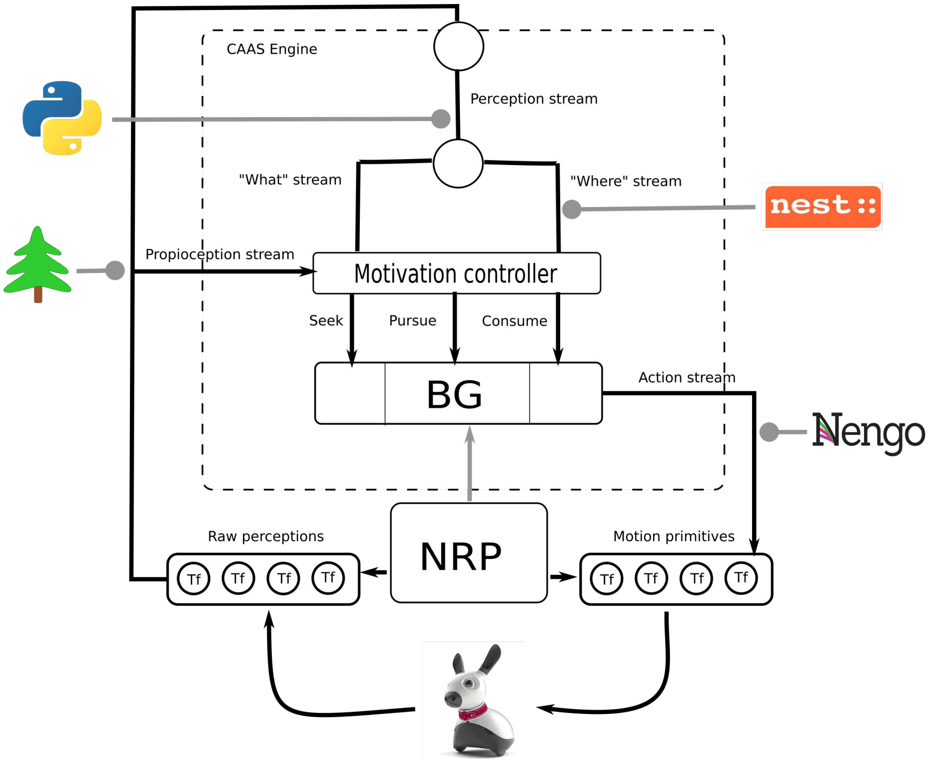

Fig. 1. CaaS scaffold architecture. The robot is modeled as a set of transfer functions that are fed the raw sensory streams and consume the action events. Each stream is fed by events generated by different models created in different simulators (for brain models, these might include models based on PyNN, NEST, Nengo, etc). The set of streams shown in this particular figure is chosen as the minimal required to generate autonomous behaviour in MiRo.

time frontier) for the models to consume [27]. This acts as a very fine grained long term memory that can be processed and aggregated offline for future use depending upon the particular application needs.

4. The particular motor commands of the robot are consumers of the action stream as well; they control different aspects of the robot independently to develop approaching and avoiding behaviours, or specific gestures.
5. The Neurorobotics platform [9] orchestrates the execution of the CaaS along with the communications with the robot through the motor commands which are, effectively, transceiver functions.

3 Results

In order to illustrate the concepts of the previous section, we developed a toy architecture with simple behaviours for the MiRo-e animal-like robot. The robot will have a *green* motivational system that will care about green stimuli in the field of vision and a *red* one which will do the same with the red stimuli. Possible behaviours are approach green, approach red, and explore (figure 2).

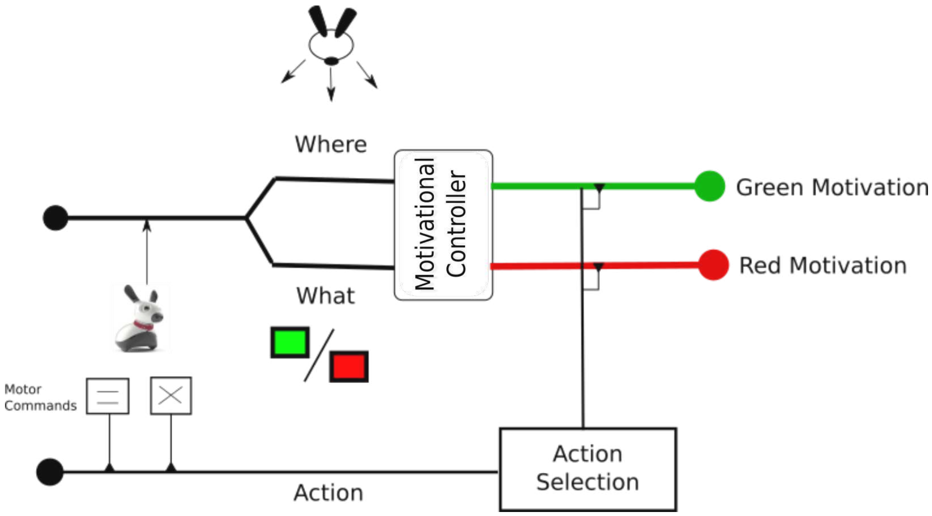

Fig. 2. Example application of the CaaS. The CaaS offers 7 streams the models can produce events into and consume events from. The raw input from the camera is consumed by basic detectors that generate *where* and *what* events. The events are consumed by simple NEST simulations that generate potential approach events onto the motivational streams. The motivational controller modulates those actions according to the internal state and generate additional *explore* commands. The final selection is performed by a basal ganglia module. The patterns shown in the motor commands correspond to two different patterns of connections in Braitenberg vehicles.

When a frame from the cameras is available, the robot generates an event in the raw input stream of the architecture. Two simple detectors consume the camera stream and produce event in the *where* and *what* stream. The *where* information is obtained by splitting MiRo's field of view in three equal parts and scanning for the appropriate colors. Two simple NEST models, implemented as in [25] take the *where* information, associated to each *what*, and generate approach commands in the corresponding motivational stream (green/red in figure 2).

Afterwards, both motivational streams are transformed by the motivational controller into two action streams corresponding to each motivation; each with either an *explore* or an *approach* command. The motivational controller itself is a switching dynamical system described elsewhere [3]. The basal ganglia selects the current action based upon the motivational state and the presence/absence of the stimuli. Finally, the commands are relayed to the transceiver functions which activate the wheels (figure 3)

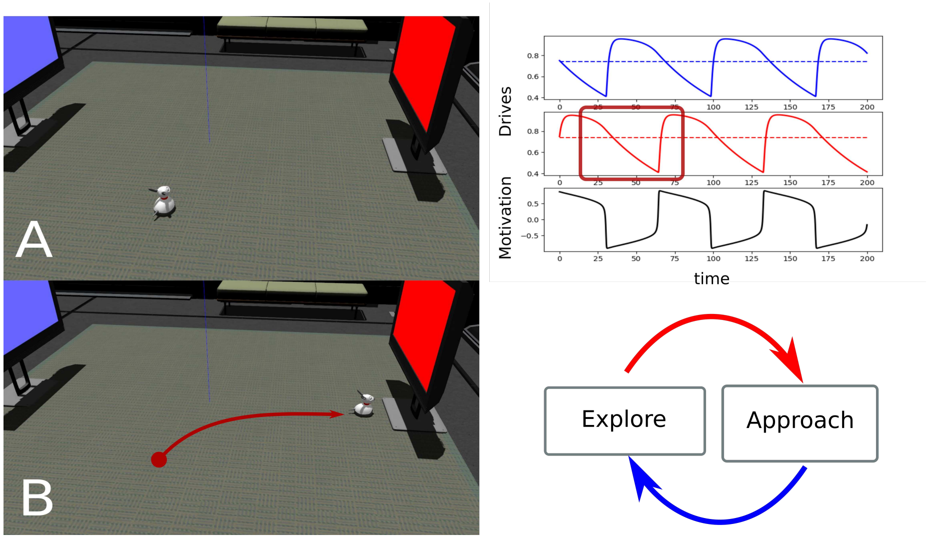

Fig. 3. Example two-motivations Braitenberg vehicle using the CaaS. The raw image from MiRo. At the start of the simulation MiRo explores until the appropriate stimulus is in the field of view. The motivational state (right) fluctuates between two motivations. When the need for red is low (square), the NEST simulation will be allowed to control the robot until the need is satisfied (being close enough to the stimulus). After that, a new cycle of exploration starts.

4 Conclusions

We have presented a loosely coupled scaffold for Cognitive architectures we call CaaS, oriented towards the development of SARs within the Human Brain Project, and have demonstrated proof-of-concept for a simple model of motivated behaviour.

This technology presents three main advantages compared to alternative solutions.

Firstly, it avoids direct communication between models by defining a scaffold of streams with specific data formats that all the models have to comply to and that have semantic significance for the overall behaviour expected from the robot. In this way, we shift the focus from the particular technical aspects of the communication between models, towards the functional role a specific model should play in the behaviour.

Second, by using Kafka technology, the design automatically logs the history of past events that have happened to the robot, or to the environment, and makes this available for the different models to consume, allowing greater flexibility and increasing the potential for applications in which long-term analyses of event are necessary. Third, even when some models take longer that the time simulation time-step to deliver particular events, the architecture guaranties behaviour at each step and can be augmented by reflex mechanisms that will mitigate for undesirable delays.A potential disadvantage of this approach are its storage and computation requirements, however, we would advocate for a mixed solution in which both embedded systems and large (off-platform) databases could be employed.

Finally, this works's contribution to the HBP's Neurorobotics platform (NRP) allows modellers to create high-level cognitive architectures by combining lower-level models of brain regions and functions in a manner that is largely agnostic concerning specific model implementations or levels of description.

5 Future work

As a further test case for the CaaS, and its digital twin potential, we plan to implement a modularised version of a more substantial cognitive architecture for the MiRo-e robot that integrates emotional expression and the capacity to align emotional state with that of the user (see [26]). MiRo-e is currently being evaluated as a potential therapy tool, with children and older adults, by the authors. We therefore hope that this extended brain-based cognitive architecture could be usefully applied for applications in hospitals, care homes, and other institutions where SARs could be deployed as therapy or emotional-support tools. A simulated version of MiRo-e has already been integrated with the NRP and will allow the usefulness of the device to be explored and assessed easily, freely, and without access to a physical robot, as well as easing further development of the robot's functionality. The NRP thus already functions as a robotic digital twin.

Acknowledgements This work was supported by the EU H2020 Programme as part of the Human Brain Project (HBP-SGA3), and specifically, through the CATRA (Cognitive Architecture for Therapy Robots and Avatars) project which was supported by the EBRAINS Research Infrastructure Voucher Programme. We thank the anonymous reviewers for the useful comments on earlier versions of the paper.

Declaration of Interest TJP is a director and shareholder of Consequential Robotics Ltd which develops the MiRo-e robot. The other authors have no competing interests.

References

1. Amoretti, M., Reggiani, M. (2010). Architectural paradigms for robotics applications. Advanced Engineering Informatics, 24(1), 4-13.
2. Prescott, T. J., Redgrave, P., Gurney, K. (1999). Layered control architectures in robots and vertebrates. Adaptive Behavior, 7(1), 99-127.
3. Jimenez-Rodriguez, A., Prescott, T. J., Schmidt, R., and Wilson, S. (2020, July). A framework for resolving motivational conflict via attractor dynamics. In Conference on Biomimetic and Biohybrid Systems (pp. 192-203). Springer, Cham.
4. Mitchinson, B., Chan, T. S., Chambers, J., Pearson, M., Humphries, M., Fox, C., ... and Prescott, T. J. (2010). BRAHMS: Novel middleware for integrated systems computation. Advanced Engineering Informatics, 24(1), 49-61.
5. Arkin RC, Fujita M, Takagi T, et al. (2003) An ethological and emotional basis for human–robot interaction. Robotics and Autonomous Systems 42(3): 191-201.
6. Asada M, MacDorman KF, Ishiguro H, et al. (2001) Cognitive developmental robotics as a new paradigm for the design of humanoid robots. Robotics and Autonomous Systems 37(2): 185-193.
7. Cangelosi A, Schlesinger M and Smith LB (2015) Developmental robotics: From babies to robots. Cambridge, MA: MIT Press.
8. Cross ES, Hortensius R and Wykowska A (2019) From social brains to social robots: applying neurocognitive insights to human–robot interaction. Philosophical Transactions of the Royal Society B: Biological Sciences 374(1771): 20180024.
9. Falotico, E., Vannucci, L., Ambrosano, A., Albanese, U., Ulbrich, S., Vasquez Tieck, J. C., ... and Gewaltig, M. O. (2017). Connecting artificial brains to robots in a comprehensive simulation framework: the neurorobotics platform. Frontiers in neurorobotics, 11, 2.
10. Kabacińska K, Prescott TJ and Robillard JM (2020) Socially Assistive Robots as Mental Health Interventions for Children: A Scoping Review. International Journal of Social Robotics. DOI: 10.1007/s12369-020-00679-0.
11. Kachouie R, Sedighadeli S and Abkenar AB (2017) The Role of Socially Assistive Robots in Elderly Wellbeing: A Systematic Review. In: Rau P-LP (ed) Cross-Cultural Design 9th International Conference.
12. Kreps, J., Narkhede, N., and Rao, J. (2011, June). Kafka: A distributed messaging system for log processing. In Proceedings of the NetDB (Vol. 11, pp. 1-7).
13. Kotseruba I, Gonzalez OJA and Tsotsos JK (2016) A Review of 40 Years of Cognitive Architecture Research: Focus on Perception, Attention, Learning and Applications. ArXiv abs/1610.08602.

14. Logan DE, Breazeal C, Goodwin MS, et al. (2019) Social Robots for Hospitalized Children. Pediatrics 144(1): e20181511.
15. Lungarella M, Metta G, Pfeifer R, et al. (2003) Developmental robotics: a survey. Connection Science 15(4): 151-190.
16. Mann JA, MacDonald BA, Kuo IH, et al. (2015) People respond better to robots than computer tablets delivering healthcare instructions. Computers in Human Behavior 43: 112-117.
17. Marti P, Bacigalupo M, Giusti L, et al. (2006) Socially Assistive Robotics in the Treatment of Behavioural and Psychological Symptoms of Dementia. The First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics, 2006. BioRob 2006., 483-488.
18. Mitchinson B and Prescott TJ (2016) MIRO: A Robot “Mammal” with a Biomimetic Brain-Based Control System. In: 5th International Conference of Biomimetic and Biohybrid Systems (eds Lepora NF, Mura A, Mangan M, et al.), Edinburgh, UK, pp.179-191. Switzerland: Springer International Publishing.
19. Prescott TJ and Robillard JM (2021) Are friends electric? The benefits and risks of human-robot relationships. iScience, 24, 101993.
20. Ostrowski AK, DiPaola D, Partridge E, et al. (2019) Older Adults Living With Social Robots: Promoting Social Connectedness in Long-Term Communities. IEEE Robotics and Automation Magazine 26(2): 59-70.
21. Perugia G, Díaz-Boladeras M, Català-Mallofré A, et al. (2020) ENGAGE-DEM: A Model of Engagement of People with Dementia. IEEE Transactions on Affective Computing. DOI: 10.1109/TAFFC.2020.2980275. 1-1.
22. Pu L, Moyle W, Jones C, et al. (2018) The Effectiveness of Social Robots for Older Adults: A Systematic Review and Meta-Analysis of Randomized Controlled Studies. The Gerontologist 59(1): e37-e51.
23. Shibata T and Wada K (2011) Robot Therapy: A New Approach for Mental Healthcare of the Elderly – A Mini-Review. Gerontology 57(4): 378-386.
24. Verschure PFMJ (2012) The Distributed Adaptive Control Architecture of the Mind, Brain, Body Nexus. Biologically Inspired Cognitive Architecture - BICA 1(1).
25. Neurorobotics.net. 2022. Tutorial setup — HBP Neurorobotics Platform 3.2.0 documentation. [online] Available at: https://neurorobotics.net/Documentation/latest/nrp/modules/CLE/hbp_nrp_cle/tutorials/setup.html [Accessed 4 July 2022].
26. Robillard JM and Hoey J (2018) Emotion and Motivation in Cognitive Assistive Technologies for Dementia. Computer, 51(3), 24–34.
27. Apache Kafka. 2022. Apache Kafka. [online] Available at: <https://kafka.apache.org> [Accessed 4 July 2022].