A Hybrid Multi-Layer Architecture for Autonomous Vehicles Utilising a Hierarchical Perception-Action Dream Simulation Mechanism

David Windridge, Iain A. D. Gunn
Department of Computer Science
Middlesex University
London
{d.windridge,i.gunn}@mdx.ac.uk

Mauro Da Lio, Alessandro Mazzalai
Department of Industrial Engineering
University of Trento
Trento, Italy
{mauro.dalio, alessandro.mazzalai}@unitn.it

Serge Thill, Henrik Svensson
School of Informatics
University of Skövde
Skövde, Sweden
{serge.thill, henrik.svensson}@his.se

Sean Anderson, Alexander Blenkinsop, Kevin Gurney, Sebastian James
Department of Automatic Control and Systems Engineering
Univeristy of Sheffield
Sheffield, UK
{s.anderson, a.blenkinsop, k.gurney, seb.james}@sheffield.ac.uk

Rafael Math, Elmar Berghöfer, Mehmed Yueksel
Deutsche Forschungszentrum für Künstliche Intelligenz GmbH
Saarbrücken, Germany
{rafael.math, elmar.berghoefer, mehmed.yueksel}@dfki.de

Andrea Saroldi, Luisa Andreone
Centro Ricerche Fiat S.C.p.A.
Orbassano (TO), Italy
andrea.saroldi@crf.it

Hermann-Josef Heich
Heich Consult GmbH
Hürth, Germany
heich@heich-consult.de

***Abstract*—Deep neural architectures have been proposed as a solution for Autonomous Vehicles. However, deep learners require very large sets of training data such that meeting current requirements of autonomous vehicle safety is intrinsically hard to achieve. We here propose a hybrid multi-layer architecture, featuring a biologically inspired separation between the tasks of action priming and action selection, that extends the principle of hierarchical Perception-Action learning via a dream simulation mechanism to greatly extend the utility of training data during learning.**

***Index Terms*—Perception-Action Learning, Automated driving; Co-Driver Agent; Artificial Cognitive Systems; Learning by simulations; Simulation Hypothesis of Cognition.**

I. Introduction

A human-driven car has an expected average fatality rate of 1.08 (fatal accidents) per 100 million miles average over all driving situations [1]. Automated driving systems will be required to improve on this to, say, 1 billion miles guaranteed fatality-free driving without human supervision, such as at level 3 or grater of the SAE automation level scale [2]. Deep Machine learning approaches, e.g. [3], thus utilize very large databases to implement a single function (lane following, in the case of [3]).

To address the challenge of learning more efficiently, the European H2020 project Dreams4Cars[1] utilizes a hybrid multi-layer architecture [4] built on the principle of hierarchical Perception-Action learning (as in the more limited previous work [5]) that is capable of extending the utility of its training data via a dream simulation mechanism. The architecture covers the complete perception to action loop in a biologically plausible model that separates action priming and action selection, with parallel priming on many potential actions and subsequent adaptive selection.

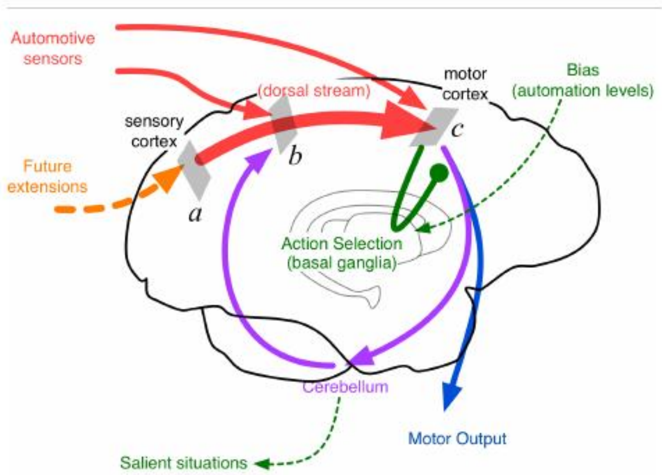

Fig. 1. Agent sensorimotor architecture

[1]www.Dreams4Cars.eu

II. ARCHITECTURE

We exploit theories of embodied and episodic simulation mechanisms [6], according to which thoughts, and dreams, are equivalent to simulated action/perception chains. We adopt a biologically inspired (see [7]) architecture with three loops:

(A) A "dorsal stream", which enacts hierarchical layered inverse models that generate affordances from the sensory input.
(B) An action selection mechanism ("basal ganglia") that operates on several levels of the hierarchy.
(C) A "cerebellum" that learns forward models.

A. Dorsal stream – layered control

The dorsal stream takes sensory input (in any available form, e.g. LIDAR, digital maps, cameras) and gradually morphs this input into a representation of possible actions (affordances) encoded as patterns of activity in the "motor cortex". Overall the dorsal stream aims at implementing simultaneous parallel action priming according to the affordance competition hypothesis by Cisek [8].

Layered organization (subsumption) in the dorsal stream allows for strategies of a higher level of complexity (complex sequences of actions). Learning and optimizing these strategies (at all levels) equates to learning the mapping from perception to affordances. The goal of the dorsal stream is thus to produce many potential actions, not just one (the selection of the actual action to be executed will be carried out later). The fact that Dreams4Cars maps perception-to-(many)-affordances and then affordances-to-(one)-action in two steps, instead of perception-to-(one)-action in a single step, differentiates Dreams4Cars from other examples of learning perception-action maps such as, e.g., [3].

Separating affordance generation from action selection achieves: 1) better adaptive behaviour, resulting from selection among a pool of potential actions; and 2) the ability for the agent to develop new behaviours (i.e. affordances) without needing to retrain the action selection mechanism (this translates to better scalability to complex situations).

Affordances are represented by active regions of the cortex and learning new motor strategies means updating the dorsal stream so that new affordances appear as new active regions. Motor cortex intensities encode the merit (inverse cost) of each action, or its *saliency* from the point of view of the agent. Actions leading to collisions will be completely inhibited; actions violating priorities of lesser importance will be only partially inhibited, so that e.g. the system will choose to break a traffic rule if this is necessary to avoid a collision (in this case, partial inhibition of traffic rules forms a shallow peak in the motor cortex).

B. Basal Ganglia (BG) – action selection

The action-selection mechanism may operate at several levels within the hierarchical subsumption architecture. At the lowest level, it operates on the motor cortex: it takes the motor cortex as input and returns a copy of the motor cortex with only one active peak - the selection occurs by suppressing all the affordance peaks except one. At a higher levels, output consists of grounded symbolic representations of salient environmental features and potential actions.

The action selection is analogous to a CNN max-pooling layer. However, in biological systems the selection mechanism is not strictly Winner Takes All (WTA); rather it is speculated that the BG implement a more sophisticated algorithm such as MSPRT (multi-hypothesis sequential probability ratio test) [9] which carries out optimal decision making between alternative actions with time and error rate constraints. The action-selection mechanism in Dreams4Cars is inspired by the broader biological principles.

C. Cerebellum – forward models

Forward models should produce an anticipation of the entire sensory input. We may break down this function into the prediction of the host vehicle trajectory, which depends on the action that is selected (or potentially selected) by the agent, and the prediction of the other agents behaviours.

Dreams4Cars will evaluate two approaches:

Hybrid analytical-learning: learning the parameters of a vehicle model and supplementing it with learning the un-modelled aspects via a general learning framework utilising a-priori knowledge constraining the learning process. An important reason for a hybrid approach, using locally weighted projection regression (LWPR) for the learnt part, is that symbolic derivatives for the forward dynamics are available (both the model equations and LWPR allows symbolic differentiation so that a Variational formulation of Optimal Control can be used in connection with logic-base reasoning.

Neural Networks. Deep Neural Networks may be used for either learning the un-modelled dynamics (replacing LWPR) or for learning the complete dynamics. Also, in this case we may exploit some ideas related to inversion of DNN such as to generate the expected input from patterns of activation representing various output symbols [10] (interpolation between symbols becomes possible and generates interpolated input).

In biological systems forward models have several uses: a) online, they may be used to produce (overt) expectations of the sensory feedback, and to enhance and process sensory information; b) in covert actions (offline use) they may be used for motor imagery and various forms of simulation of actions.

Dreams4Cars foresees 3 particular uses:

1) A self-monitoring system to detect anomalies in vehicle dynamics.
2) Detection of novelties (online use), i.e., mismatch between agent prediction and what happens at the higher-levels of the sensorimotor architecture to annotate salient situations for future dreams.
3) Learning a model of the world for offline simulations (dreams)

III. GENERATION OF DREAMS

Within the architecture, dreams reactivate the offline control system as if it were interacting with the controlled entity.

They thus enable the biological agent to consider hypothetical situations to increase its ability to handle situations in the online state [11], [12].

The dream-state simulates previous and novel events; it can utilise goal directed scenarios to explore alternate paths. However, one mode in particular enables very efficient use of training data: simulations that recombine aspects of previously encountered events into new events. Novel events can be achieved by rearranging percepts, including: instance changing (e.g. change of vehicle type in a given situation), modifying trajectories, change of environment, rotating forward emulators. One research direction is experimenting with convolutional-decovovolutinal networks which somewhat mimic the convergence-divergence neural architecture posited by Meyer and Damasio [13]. In particular the future space in this type of DNN may be the common points for representation of past experiences (by compression using the coder part of the net), recall of memorised situations (by expansion using the decoder part to an associated autoencoder), and motor planning (by expanding with the net decoder part). (Note though that subsumption by higher levels in our architecture must be respected when plannning motor activity.)

In this context, top-down exploratory instantiations of the subsumption architecture can be used to generate goal directed simulations. Randomized selection of perceptual goals within the perception-action hierarchy amounts to motor babbling in a manner analogous to the learning process of infant humans. In particular, this motor babbling is carried out top-down; higher-level percepts thus become the goal states of actions parameterized at each hierarchical level in a subsumptive fashion [14], [15], [16] . This mechanism is thus sufficiently flexible to allow for first-order symbolic logical reasoning processes (such as in relation to the Highway code) to be integrated with, e.g. optimal control mechanisms on lower levels (OC thus acts to optimize motor primitives within this context). Dream generation is thus enacted by top-down logical variable instantiation in a manner analogous to generative deep-learning approaches.

References

[1] NHTSA, "Fatality analysis reporting system (FARS) encyclopedia," https://www-fars.nhtsa.dot.gov/Main/index.aspx.

[2] SAE, "Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles," http://standards.sae.org/j3016_201609.

[3] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. Jackel, M. Monfort, U. Muller, J. Zhang, and X. Zhang, "End to end learning for self-driving cars," 2016, arXiv preprint arXiv:1604.07316.

[4] Dreams4Cars, "Deliverable D1.2, system architecture," http://www.dreams4cars.eu/en/d/project/deliverables.

[5] M. D. Lio, F. Biral, E. Bertolazzi, M. Galvani, P. Bosetti, D. Windridge, A. Saroldi, and F. Tango, "Artificial co-drivers as a universal enabling technology for future intelligent vehicles and transportation systems," *IEEE Trans. on Intelligent Transportation Systems*, vol. 16, no. 1, 2015.

[6] H. Svensson and S. Thill, "Beyond bodily anticipation: internal simulations in social interaction." *Cognitive Systems Research*, vol. 40, pp. 161–171, 2016.

[7] S. Thill, D. Caligiore, A. M. Borghi, T. Ziemke, and G. Baldassarre, "Theories and computational models of affordance and mirror systems: An integrative review," *Neuroscience & Biobehavioral Reviews*, vol. 37, no. 3, pp. 491–521, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0149763413000134

[8] P. Cisek, "Cortical mechanisms of action selection: the affordance competition hypothesis," *Philos. Trans. R. Soc. B Biol. Sci.*, vol. 362, no. 1485, pp. 1585–1599, 2007.

[9] R. Bogacz and K. Gurney, "The basal ganglia and cortex implement optimal decision making between alternative actions," *Neural computation*, vol. 19, no. 2, pp. 442–477, 2007.

[10] A. Mahendran and A. Vedaldi, "Understanding deep image representations by inverting them," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[11] R. Chrisley, "Cognitive map construction and use: A parallel distributed processing approach," in *Connectionist models: Proceedings of the 1990 summer school*, 1990.

[12] J. Tani and S. Nolfi, "Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems," *Neural Netw.*, vol. 12, no. 7–8, pp. 1131–1141, 1999.

[13] K. Meyer and A. Damasio, "Convergence and divergence in a neural architecture for recognition and memory," *Trends in Neurosciences*, vol. 32, no. 7, pp. 376–382, Jul. 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0166223609000903

[14] D. Windridge, "Emergent intentionality in perception-action subsumption hierarchies," *Frontiers in Robotics and AI*, vol. 4, p. 38, 2017.

[15] D. Windridge and J. Kittler, "Perception-action learning as an epistemologically-consistent model for self-updating cognitive representation," *Adv. Exp. Med. Biol.*, vol. 657, pp. 95–134, 2010.

[16] D. Windridge, M. Felsberg, and A. Shaukat, "A framework for hierarchical perception-action learning utilizing fuzzy reasoning," *IEEE Trans. Syst. Man Cybern. Part B Cybern.*, vol. 43, no. 1, pp. 155–169, 2012.