Dream-like simulation abilities for automated cars

DREAMS4CARS

Grant Agreement No. 731593

Deliverable:	D7.3. – Final Report (public summary)
Dissemination level:	PU - Public
Delivery date:	17/December/2019
Status:	Final

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731593

<table>
<tr><td>Deliverable Title</td><td colspan="4">Final Report (public summary)</td></tr>
<tr><td>WP number and title</td><td colspan="4">WP7 - Management</td></tr>
<tr><td>Lead Editor</td><td colspan="4">Mauro Da Lio, UNITN</td></tr>
<tr><td rowspan="3">Contributors</td><td colspan="4">Henrik Svensson, HIS</td></tr>
<tr><td colspan="4">Andrea Saroldi, CRF</td></tr>
<tr><td colspan="4"></td></tr>
<tr><td>Creation Date</td><td colspan="2">7/December/2019</td><td>Version number</td><td>0.5</td></tr>
<tr><td>Deliverable Due Date</td><td colspan="2">31/December/2019</td><td>Actual Delivery Date</td><td>17/December/2019</td></tr>
<tr><td rowspan="4">Nature of deliverable</td><td>X</td><td colspan="3">R - Report</td></tr>
<tr><td></td><td colspan="3">DEM – Demonstrator, pilot, prototype, plan designs</td></tr>
<tr><td></td><td colspan="3">DEC – Websites, patents filing, press&media actions</td></tr>
<tr><td></td><td colspan="3">O – Other – Software, technical diagram</td></tr>
<tr><td rowspan="3">Dissemination Level/ Audience</td><td></td><td colspan="3">PU – Public, fully open</td></tr>
<tr><td>X</td><td colspan="3">CO - Confidential, restricted under conditions set out in MGA</td></tr>
<tr><td></td><td colspan="3">CI – Classified, information as referred to in Commission Decision 2001/844/EC</td></tr>
</table>

Version	Date	Modified by	Comments
0.1	7/December/2019	Mauro Da Lio	Initial version.
0.2	16/December/2019	Mauro Da Lio	Version for review.
0.3	16/December/2019	Andrea Saroldi	Peer Review comments.
0.4	16/December/2019	Henrik Svensson	Peer Review comments.
0.5	17/December/2019	Mauro Da Lio	Final version.

1 Final Report Public Summary

Deliverable D7.2 (the confidential counterpart of this document) gives a retrospective analysis of how our vision of Cognition for Autonomous Driving evolved.

Overall the present document can be regarded as summary of that analysis. It is also an anticipation of a position paper that will be shortly submitted (so details that are not disclosed here are saved for a future journal work).

For what is relevant for this document, in our analysis, we focus on the challenges and opportunities for Cognition in Autonomous Driving.

2 Challenges in Autonomous Driving

We start observing that the assumptions that humans are bad drivers and that driving can be "automated" are superficial. In reality the average fatality rate of human driving is 1 every 100 million miles. This figure includes all types of drivers (inattentive, young, drunk, etc.). A senior attentive driver would be much better. Hence, if safety has to be used as the main motivation for automated vehicles, mature automated vehicles should target much better figures [1], [2]. On the other hand, reports say that the progress in the number of autonomously driven miles is not very fast [3], [4] and plagued by many corner cases, e.g.: [5].

An example (outlined in D7.2) of a not-so-rare motorway situation that could become a corner case is used to illustrate the difference between a Deep Neural Network for semantic segmentation (often referred as being full-fledged AI in itself, which is not true) and Cognition; only in the latter case, an intelligent system can autonomously make *predictions* of what might happen and use predictions to elaborate *behaviours*.

The architecture almost universally adopted for automated driving systems is the *sense-think-act* paradigm [6], [7]. We criticize this choice (which is not a credible model for human intelligence [8]–[21]) arguing that it is this architecture that carries in its own nature the reasons for the endless emergence of corner cases (see also [22]). It actually works well in situations corresponding to symbols and behaviours that have been programmed, but it is not economical when the number of corner cases (each with different graduation) is virtually infinite. A brute force approach may suggest that, with sufficient testing and developments, all corner cases will be sooner or later discovered, programmed, and a final (how much complex?) *complete* software will be developed; but there actually is little evidence of this assumption (e.g., [3], [4]). In addition, regardless of whether, given enough resources, sooner or later a brute force approach will yield us a fully functioning Level 4-5 automated car, the question remains that this way is very expensive in terms of resources to be committed and yields very complex software systems.

End-to-end approaches are also considered [23]. Even if it is possible to explain to which features end-to-end trained networks respond [24], that still questions how decisions are taken. Besides that –not to mention example cases of possible failures [25]–, a monolithic end-to-end network does not produce behaviours as a competition between affordances: the priming of the possible actions and the selection of the best option are intimately entangled (there is no indication of which other potential decisions might have been considered and discarded). Furthermore, the network is not adaptive: the output of the network is not broken into hierarchical levels of intentions. So, one thing might be the intention of achieving one particular position in the lane, another is the exact steering angle to be used for that, which depends on the particular dynamics of the driven vehicle. If the vehicle were different, or if the vehicle dynamics changes because of environmental and operating conditions, the network might need to be retrained.

3 Opportunities for Cognition in Autonomous Driving

Dreams4Cars was informed by the idea that thinking is a simulation process [26]: loops of simulated actions that elicits simulated perception. There may be *inline* simulations when, for example, an agent makes simulations in parallel with the course of actions. There may also be *offline* simulations where at a separate time (with less time constraints) alternative situations are imagined and used to elaborate appropriate action strategies. This involves both high-level behaviours (e.g., how to negotiate a particular traffic situation) as well as lower-level control (e.g., which exact steering control to use on a slippery road). Making an (imperfect) parallelism with human offline forms of cognition, we could, for example, argue that in REM sleep higher level behaviours are synthetized; whereas there is some evidence that at stage 2 sleep inverse model for motor control may be synthetized [27].

Of course, during the development of Dreams4Cars many other literature findings have influenced our work. We mention, among the others, the organization of brain neural networks in the cerebellum [28], [29], in the dorsal stream [30] and in the basal ganglia [31]–[33]; the topographic organization of many brain areas and in particular of the motor cortex [34], [35]; the affordance competition hypothesis [18]; where and how predictions occur in the brain [17]; the notion of episodic versus embodied simulation [36] (also related to distinction between declarative and procedural prediction [17]); intentional biasing action selection [37] (this in particular has provided the solution for embedding traffic rules in the agent); subsumption [8] and layered control architectures [38].

The main focus of Dreams4Cars has been on offline simulations and was concentrated on two main innovations: agent architecture and offline synthesis of behaviours and control.

3.1 Agent architecture

For the agent architecture our starting point was a subsumption architecture [8] with posterior centralized action selection: namely a layered control architecture [38].

We have then improved the Agent architecture with several ideas from the above cited literature. The architecture hence goes beyond subsumption architectures and, for certain aspects, also beyond layered control architectures. Scientifically, it is an interesting example of application of many biological principles [39]–[43].

However, apart from the biological inspiration, the agent architecture presents several technical and scientific innovations, especially if compared to the current automotive methodologies (the sense-think act paradigm, not to mention the less realistic end-to-end training). Worth to be stressed is, in particular, the topographic organization of the motor space with its fall-backs among which explainability and modularity, the way in which traffic rules can be incorporated into the agent via biasing mechanisms, the minimum commitment principle, the minimum intervention principle and the principle of lower veto that makes a sand box for programmed rules.

The emergence of complex behaviours from the agent architecture and its founding principles is also important because it implies a fundamental lean and simple software, composed of a very compact core which, in turn, is economical and easy to maintain.

Biasing mechanisms, inspired by works like [37], have been exploited to integrate traffic rules (and logical module) into the agent sensorimotor behaviours.

3.2 Synthesis of control and behaviours with off line mental imagery

The process for the synthesis of behaviours is the second major contribution. It is based on the learning of models of the world that are then manipulated to synthetize action strategies. It is very efficient both with respect to produced behaviours and for sample efficiency. It bootstraps a hierarchical sensorimotor architecture where each level learning is accelerated by what was learned at the immediately precedent level. The obtained sensorimotor system replaces corresponding aspects of traditional motor control (MPC, Optimal Control, Trajectory planning, etc., [6], [44]–[46]). In addition, elements of the process can be adopted individually and this is expected to facilitate exploitation and diffusion of the technology.

While the synthesis process creates the blocks for higher-levels of simulation abstractions, these blocks become also available for the very highest-level, which may be implemented via Reinforcement Learning. This way Reinforcement Learning may proceed on a low-dimensional state-action space and be, in principle, more efficient because on top of efficient neural networks [47].

It is also worth observing that, beginning at the level of motor primitives, the main goal for learning inverse models is not just learning the action for one specific goal, but a function that returns the value of actions for posterior action selection. This value function is the salience stored in the “motor cortex”. It is functionally equivalent to the reward function in Reinforcement Learning, but its computation is not achieved via a trial and error exploration (the RL way) but via a synthesis process based on the offline manipulation of learned models of the world (the inspiring idea of Dreams4Cars).

Bibliographical References

[1] N. Kalra e S. M. Paddock, «Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?», *Transp. Res. Part Policy Pract.*, vol. 94, pagg. 182–193, 2016.

[2] *Gil Pratt, Toyota CES 2017 Live Stream*. .

[3] V. V. Dixit, S. Chand, e D. J. Nair, «Autonomous Vehicles: Disengagements, Accidents and Reaction Times», *PLOS ONE*, vol. 11, n. 12, pag. e0168054, dic. 2016.

[4] C. Lv *et al.*, «Analysis of autopilot disengagements occurring during autonomous vehicle testing», *IEEECAA J. Autom. Sin.*, vol. 5, n. 1, pagg. 58–68, gen. 2018.

[5] NTSB, «Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian». [Online]. Available at: https://www.ntsb.gov/news/events/Pages/2019-HWY18MH010-BMG.aspx. [Accessed: 06-dic-2019].

[6] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, e E. Frazzoli, «A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles», *IEEE Trans. Intell. Veh.*, vol. 1, n. 1, pagg. 33–55, mar. 2016.

[7] C. Badue *et al.*, «Self-Driving Cars: A Survey», *ArXiv190104407 Cs*, ott. 2019.

[8] R. Brooks, «A robust layered control system for a mobile robot», *IEEE J. Robot. Autom.*, vol. 2, n. 1, pagg. 14–23, 1986.

[9] T. Van Gelder, «What Might Cognition Be, If Not Computation?», *J. Philos.*, vol. 92, n. 7, pagg. 345–381, 1995.

[10] J. A. Seitz, «The bodily basis of thought», *New Ideas Psychol.*, vol. 18, n. 1, pagg. 23–40, 2000.

[11] J. Gibson, «The theory of affordances», in *The Ecological Approach To Visual Perception*, Lawrence Erlbaum Associates, 1986, pagg. 127–143.

[12] F. Keijzer, «Representation in dynamical and embodied cognition», *Cogn. Syst. Res.*, vol. 3, n. 3, pagg. 275–288, 2002.

[13] L. W. Barsalou, «Grounded cognition.», *Annu. Rev. Psychol.*, vol. 59, pagg. 617–45, 2008.

[14] R. Grush, «The emulation theory of representation: motor control, imagery, and perception.», *Behav. Brain Sci.*, vol. 27, n. 3, pagg. 377–96; discussion 396–442, giu. 2004.

[15] G. Hesslow, «Conscious thought as simulation of behaviour and perception», *Trends Cogn. Sci.*, vol. 6, n. 6, pagg. 242–247, giu. 2002.

[16] J. Decety e J. Grèzes, «The power of simulation: Imagining one's own and other's behavior», *Brain Res.*, vol. 1079, n. 1, pagg. 4–14, mar. 2006.

[17] K. L. Downing, «Predictive models in the brain», *Connect. Sci.*, vol. 21, n. 1, pagg. 39–74, mar. 2009.

[18] P. Cisek, «Cortical mechanisms of action selection: the affordance competition hypothesis», *Philos. Trans. R. Soc. B Biol. Sci.*, vol. 362, n. 1485, pagg. 1585–1599, set. 2007.

[19] M. Kiefer e F. Pulvermüller, «Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions», *Cortex*, vol. 48, n. 7, pagg. 805–825, lug. 2012.

[20] M. Jeannerod, «Neural simulation of action: a unifying mechanism for motor cognition.», *NeuroImage*, vol. 14, n. 1 Pt 2, pagg. S103–9, lug. 2001.

[21] A. Guillot e C. Collet, «Duration of Mentally Simulated Movement: A Review», *J. Mot. Behav.*, vol. 37, n. 1, pagg. 10–20, gen. 2005.

[22] M. Da Lio *et al.*, «Artificial Co-Drivers as a Universal Enabling Technology for Future Intelligent Vehicles and Transportation Systems», *IEEE Trans. Intell. Transp. Syst.*, vol. 16, n. 1, pagg. 244–263, 2015.

[23] M. Bojarski *et al.*, «End to end learning for self-driving cars», *ArXiv Prepr. ArXiv160407316*, 2016.

[24] M. Bojarski *et al.*, «Explaining How a Deep Neural Network Trained with End-to-End Learning Steers a Car», *ArXiv Prepr. ArXiv170407911*, 2017.

[25] A. Rosenfeld, R. Zemel, e J. K. Tsotsos, «The Elephant in the Room», *ArXiv180803305 Cs*, ago. 2018.

[26] G. Hesslow, «The current status of the simulation theory of cognition.», *Brain Res.*, vol. 1428, pagg. 71–9, gen. 2012.

[27] C. T. Smith, J. B. Aubrey, e K. R. Peters, «Different Roles for REM and Stage 2 Sleep In Motor Learning: A Proposed Model», pag. 24.

[28] J. Porrill, P. Dean, e S. R. Anderson, «Adaptive filters and internal models: Multilevel description of cerebellar function», *Neural Netw.*, vol. 47, pagg. 134–149, 2012.

[29] D. M. Wolpert, R. C. Miall, e M. Kawato, «Internal models in the cerebellum.», *Trends Cogn. Sci.*, vol. 2, n. 9, pagg. 338–47, set. 1998.

[30] K. Meyer e A. Damasio, «Convergence and divergence in a neural architecture for recognition and memory», *Trends Neurosci.*, vol. 32, n. 7, pagg. 376–382, lug. 2009.

[31] R. Bogacz e K. Gurney, «The basal ganglia and cortex implement optimal decision making between alternative actions», *Neural Comput.*, vol. 19, n. 2, pagg. 442–477, 2007.

[32] K. Doya, «What are the computations of the cerebellum, the basal ganglia and the cerebral cortex?», *Neural Netw.*, vol. 12, n. 7–8, pagg. 961–974, ott. 1999.

[33] P. Redgrave, T. J. Prescott, e K. N. Gurney, «The basal ganglia: a vertebrate solution to the selection problem?», *Neuroscience*, vol. 89, pagg. 1009–1023, 1999.

[34] D. Purves *et al.*, «Functional Organization of the Primary Motor Cortex», *Neurosci. 2nd Ed.*, 2001.

[35] G. H. Patel, D. K. Michael, e L. H. Snyder, «Topographic organization in the brain: Searching for general principles», *Trends Cogn. Sci.*, vol. 18, n. 7, pagg. 351–363, lug. 2014.

[36] H. Svensson e S. Thill, «Beyond bodily anticipation: Internal simulations in social interaction», *Cogn. Syst. Res.*, vol. 40, pagg. 161–171, 2016.

[37] G. Pezzulo e P. Cisek, «Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition», *Trends Cogn. Sci.*, vol. 20, n. 6, pagg. 414–424, giu. 2016.

[38] T. J. Prescott, P. Redgrave, e K. Gurney, «Layered Control Architectures in Robots and Vertebrates», *Adapt. Behav.*, vol. 7, n. 1, pagg. 99–127, gen. 1999.

[39] M. Da Lio, A. Mazzalai, K. Gurney, e A. Saroldi, «Biologically Guided Driver Modeling: the Stop Behavior of Human Car Drivers», *IEEE Trans. Intell. Transp. Syst.*, vol. in press (http://ieeexplore.ieee.org/document/8057588/).

[40] M. Da Lio, D. Bortoluzzi, e G. P. Rosati Papini, «Modelling longitudinal vehicle dynamics with neural networks», *Veh. Syst. Dyn.*, pagg. 1–19, lug. 2019.

[41] S. James, A. Sean, e D. L. Mauro, «Longitudinal Vehicle Dynamics: A Comparison of Physical and Data-Driven Models Under Large-Scale Real-World Driving Conditions», *Veh. Syst. Dyn.*, vol. (submitted).

[42] R. Donà, G. P. Rosati Papini, M. Da Lio, e L. Zaccarian, «On the Stability and Robustness of Hierarchical Vehicle Lateral Control With Inverse/Forward Dynamics Quasi-Cancellation», *IEEE Trans. Veh. Technol.*, vol. 68, n. 11, pagg. 10559–10570, nov. 2019.

[43] S. James e S. R. Anderson, «Linear System Identification of Longitudinal Vehicle Dynamics Versus Nonlinear Physical Modelling», in *2018 UKACC 12th International Conference on Control (CONTROL)*, Sheffield, 2018, pagg. 146–151.

[44] A. Bemporad e M. Morari, «Robust model predictive control: A survey», in *Robustness in identification and control*, vol. 245, A. Garulli e A. Tesi, A c. di London: Springer London, 1999, pagg. 207–226.

[45] E. Kim, J. Kim, e M. Sunwoo, «Model predictive control strategy for smooth path tracking of autonomous vehicles with steering actuator dynamics», *Int. J. Automot. Technol.*, vol. 15, n. 7, pagg. 1155–1164, dic. 2014.

[46] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, e D. Hrovat, «Linear time-varying model predictive control and its application to active steering systems: Stability analysis and experimental validation», *Int. J. Robust*

Nonlinear Control, vol. 18, n. 8, pagg. 862–875, mag. 2008.

[47] M. Botvinick, S. Ritter, J. X. Wang, Z. Kurth-Nelson, C. Blundell, e D. Hassabis, «Reinforcement Learning, Fast and Slow», *Trends Cogn. Sci.*, vol. 23, n. 5, pagg. 408–422, mag. 2019.