Autonomous Vehicle Architecture Inspired by the Neurocognition of Human Driving

Mauro Da Lio[1], Alice Plebe[2], Daniele Bortoluzzi[1], Gastone Pietro Rosati Papini[1] and Riccardo Donà[1]
[1]*Department of Industrial Engineering, University of Trento, Trento, Italy*
[2]*Department of Information Engineering and Computer Science, University of Trento, Trento, Italy*

Keywords: Automated Driving, Cognitive Systems.

Abstract: The realization of Autonomous vehicles is recognized as a relevant objective for the modern society and constitutes a challenge which in the last decade is concentrating a growing interest, involving both manufacturers and research institutes. The standard approach to the realization of automated driving agents is based on a well-known paradigm, consisting of the sense-think-act scheme. Even though this implements an understandable and agreeable logic, a driving agent based on such an approach needs to be tested and qualified at a level of reliability which requires a huge experimental campaign. In this position paper the scope of the problem of automated driving is widened into the cognitive sciences, where the inspiration is taken to reformulate the underlying paradigm of the automated agent architecture. In the framework of the EU Horizon 2020 Dreams4Cars Research and Innovation Action project the challenge is to design and train an automated driving agent which mimics the known human cognitive architecture and as such is able to learn from significant situations encountered (either simulated or experienced), rather than simply applying a set of fixed rules.

1 INTRODUCTION

The recent years witnessed an exponential increase in the number of industrial –and research– initiatives aimed at the implementation of automated driving. However, after the initial enthusiasms, one challenge in particular is emerging: how to prove that the engineered vehicles are safer than humans. One recent study (Kalra and Paddock, 2016) demonstrates that extensive road testing would require driving for billions of miles –which is not realistic– to provide *statistically significant* proofs that the software is safer than humans. By comparison, companies testing Autonomous Driving (AD) vehicles and reporting to the California State have accrued less than 0.5 million miles between 9/2014 and 11/2015 (Dixit et al., 2016). The problem stems from the fact that humans –contrary to superficial perception– are very reliable at driving: in the US there are about 33,000 fatalities and 2.3 million injuries per year. As large as these figures may look like, when divided by the total traveled miles (3 trillions) they correspond to very low rates of 1.09 fatalities and 77 injuries per 100,000,000 miles.

This position paper proposes a paradigm shift for the *architecture* of the driving agents. Today, the software for AD follows the *sense-think-act* paradigm, which means that the perception system produces a symbolic representation of the environment and that there is software determining the agent behavior based on this representation (a.k.a. Cartesian rationalism). This approach almost necessarily means that the software must be entirely developed and tested by human designers. Simulations may be used to reduce the amount of testing but the tested scenarios, again, have to be known. For the sense-think-act architecture, the goal of designing and testing a system capable of error-free behaviors for billions of miles is definitely a challenge.

Overall, designing a system capable of correct behavior for billions of miles means designing a system capable of acting in *autonomy*, i.e., in situations that might have not been expressly considered at the design stage. Humans have the ability to transfer knowledge of the few generic situations that they faced in a few tenths of training hours at driving schools to deal with novel situations found in a lifetime of driving. Moreover they continuously learn: for example, every critical situation (even only potential hazards) is detected and *imagery mechanisms* to study that situations at both dreaming and wake states exist, that allow humans to learn new behavioral strategies for situations that they have only imagined. That is why

Lio, M., Plebe, A., Bortoluzzi, D., Papini, G. and Donà, R.
Autonomous Vehicle Architecture Inspired by the Neurocognition of Human Driving.
In *Proceedings of the 4th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS 2018)*, pages 507-513
ISBN: 978-989-758-293-6

young licensees have far more risk than senior drivers.

Human accidents are largely due to distractions, tiredness, risk taking or driving under the effects of alcohol or drugs, aspects that will not affect artificial systems. Hence we want to take inspirations from the mechanisms underlying human autonomy and propose an artificial cognitive architecture capable of similar autonomy, and capable of increasing its abilities like senior drivers. For this, we analyze in the rest of the paper the human sensorimotor system and show how it might be taken as a model for the architecture of driving agents.

This paper is one output of the EU Horizon 2020 "Dreams4Cars" Research and Innovation action (www.dreams4cars.eu).

2 THE NEUROCOGNITIVE UNDERPINNING OF DRIVING

The ability to drive is one the many highly specialized human sensorimotor behaviors. In addition to the universal ability to walk, run, and jump, specializations that share with driving complex motor coordination include playing piano, skying, playing tennis, and the like (Wolpert et al., 2011). The history of robotics has proved how difficult it is to implement perception-action systems that lead to performances similar to human ones, even for the basic walking abilities. Automated driving may appear easier, since car motion is more prone to direct computer control than, for example, robots with human-like gait. However, we deem that the sophisticated control system that the human brain develops, when learning to drive by commanding the ordinary car interfaces, steering wheel and pedals, may reveal precious insights on how to implement a robust automatic driving system.

The task to drive is solved by the human brain with the same kind of strategy that we adopt for every sort of motor planning that requires continuous and complex perceptual feedback. For sure, how this general strategy works is far from being fully understood, and there are several competing theories, lacking a unanimous consensus. Nevertheless, a huge body of research in neuroscience and cognitive neuroscience has been produced in the past decades, that allows us to grasp some general cues useful for designing artificial car control systems.

2.1 Simulation and Emulation

Primarily, we consider approaches that belong to the simulative theories, in the account reviewed in (Hesslow, 2012). In cognition, the idea of simulation has been traditionally related with the hypothesis that our social cognition is based on "simulating" internally what is going on in the mind of others, adopting a sort of psychological theory of how (ours and other's) minds work (Gallagher, 2007). The simulation theory supported by Hesslow points to a different direction, he argues that thinking in general is explicated by simulating perceptions and actions involved in the thought, without the need of actually executing the actions, or perceiving online what is imagined. This is also the course of action our project intends to take forward (Da Lio et al., 2017). In the view of Hesslow, simulation is a general principle of cognition, explicated in at least three different components: actions, perception, and anticipation.

One of the earlier proposals in this direction is the theory of neural emulators (Grush, 2004), that bridges a close link between the engineering domain of control theory and signal processing to neural representations. According to the emulation theory, the brain is able to constructs neural models of the body and of the environment, in addition to simply engaging with them. During overt sensorimotor engagement, these models are driven by efference copies in parallel with the body and the environment, so to generate expectations of the sensory feedback. Later on, these models can also be run offline, in order to predict outcomes of different actions, and evaluate and develop motor plans. In terms of control theory, Grush's emulators are essentially forward models, and so will be called in the rest of the paper. Grush further argued that forward models can be realized by Kalman filters, pointing to models that seem to support his hypothesis (Wolpert and Kawato, 1998), and neuroscientific evidences for forward models in the cerebellum (Wolpert et al., 1998; Jeannerod and Frak, 1999). We will not use the Kalman filter hypothesis here, while we embrace the characterization of the cerebellum as site of forward models predicting the sensory effects of movements.

2.2 Imagery and Dreaming

A second tradition of research that converges into Hesslow's simulation account is perceptual imagery, especially of visual modality. Mental imagery is the phenomenon where a representation of the type created during the initial phases of perception is present, but the stimulus is not actually being perceived; such representations preserve the perceptible properties of the stimulus (Moulton and Kosslyn, 2009). Mental imagery has connections with the phenomenon of dreaming, and it has been argued (Thill and Svensson, 2011) that the function itself of dreaming, in infants through to early childhood, is to trigger and exercise the capacity of simulating.

What is the mechanism that allows primary sensorial areas to be activated both by online stimuli or by perceptual imagery remains to be unveiled. A prominent proposal is formulated in term of convergence-divergence zones (CDZs) (Meyer and Damasio, 2009). CDZs receive convergent projections from the early sensorimotor sites and send back divergent projections to the same sites. This arrangement has the first purpose to record the combinatorial organization of the knowledge fragments coded in the early cortices, together with the coding of how those fragments must be combined to represent an object comprehensively. CDZ records are built through experience, by interacting with objects. The CDZ framework can explain perceptual imagery, as it proposes that similar neural networks are activated when objects or events are processed in perceptual terms and when they are recalled from memory. This project will experiment with implementations loosely inspired by CDZs.

2.3 Main Neural Components of Simulation

The simulation theory as spelled out by Hesslow covers many aspects of human cognition, but does not address how the simulation framework might be actually realized in the brain. This crucial aspect is, instead, central in the investigations done by Cisek (Cisek, 2007; Cisek and Kalaska, 2010). In his general framework, simulation is one component of the mechanisms by which the brain selects actions and specifies the parameters or metrics of those actions. Cisek's theory is named *affordance competition hypothesis*, in which "affordance" is the term originally used by Gibson (Gibson, 1979) to refer to the action possibilities of the environment that are available to an animal. For example, for most human-beings a chair "affords" sitting, a glass "affords" grasping; and we might say that for a driver a lane "affords" either lane change or lane following and a yield line "affords" either stopping or crossing. In Cisek, "affordances" are the internal representations of the potential actions which are in constant competition for deciding the next behavior. In this project, the term affordances will be used in exactly the same sense.

The first neural component of the affordance competition hypothesis, in the context of visually guided actions, is the occipito-parietal *dorsal stream*. In the traditional division of the visual processing path into the *dorsal stream* and the occipito-temporal *ventral stream* the former builds a representation of *where* things are and the latter of *what* things are (Ungerleider and Mishkin, 1982). A more recent account of the dorsal stream proposes that its role is to mediate various visually guided actions by several substreams. For example, the lateralintraparietal (LIP) area specify potential saccade targets; the medial intraparietal (MIP) area specify possible directions for arm reaching.

The next fundamental component is made up by the basal ganglia in connection with the dorsolateral prefrontal cortex (DLPFC), and performs action selection. Since action selection is a fundamental problem faced by even the most primitive of vertebrates, it is consistent with the involvement of an ancient structure conserved throughout evolution, like basal ganglia. DLPFC appears to play the role of collecting "votes" for categorically selecting one action over others. There is a wide literature on the detailed mechanisms by which basal ganglia and prefrontal cortex interact in taking decisions (Redgrave et al., 1999; Doya, 1999; Bogacz and Gurney, 2007; Lewis et al., 2011).

The last fundamental component is the cerebellum, where an internal predictive feedback is generated, once the final selected action is released. As mentioned above, the cerebellum was also taken into account by Grush for its theory of emulators.

3 AGENT SENSORIMOTOR ARCHITECTURE

The overall architecture of the agent to be developed in this project is broadly derived from the neural components and brain strategies that allow humans to drive, and were described in the previous sections. For this reason, the main scheme of the agent, shown in Fig. 1, is overlaid on a sketched brain, and is an adaptation to the task of driving by an artificial agent of the action-selection scheme in (Cisek and Kalaska, 2010, Fig.1 p.278).

In the remaining of this section the three main processing pathways shown Fig. 1 will be described.

3.1 Inverse Model

The pathway depicted in red in Fig. 1 performs essentially the transformation of an inverse model, taking as input sensory data and producing outputs in the format of motor controls. This is the role played by the visual dorsal stream in the account of Cisek, and we adopt his terminology by calling *affordances* the final outputs of this processing pathway, (c) in the figure.

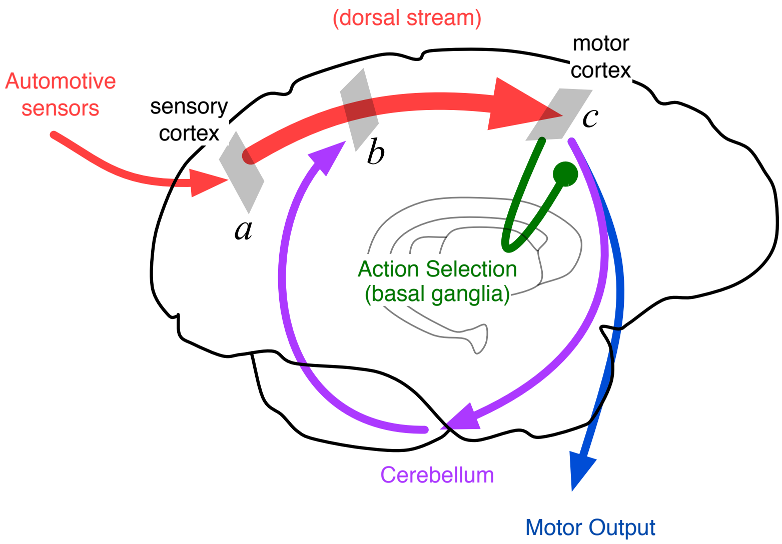

Figure 1: Main scheme of the agent architecture, the main pathways are depicted in red (dorsal stream), green (action selection) and violet (forward model); (a), (b), and (c) are progressive stages of representations in a scale from sensorial to affordances.

The format in (a) is basically in the topographic representation of the perceptual space, fusing information from sensors like camera and LIDAR. The format in (b) will be in a topology more oriented toward the egocentric control requirements. This topology will be in two dimensions, since the control space for driving can be reduced to the instantaneous lateral and longitudinal controls. The format in (c) is still in a two dimensional topology, but in terms of discrete affordances, such as “lane change”, “car follow”, and similar. As visible in Fig. 1, the architecture is arranged to receive also input from intelligent sensors that can provide data at higher levels of format, in (b) and even in (c).

This pathway, much like the biological dorsal stream, is in fact a collection of several specialized substreams, at least the following can be identified:

- road geometry;
- obstacles;
- traffic lights or other traffic directives.

There is a difference in kind among the outputs of those substreams: road geometry provides the active regions in the space of longitudinal and lateral controls, while obstacles and traffic lights compute inhibited regions.

At the core of the transformations from (a) to (c) there will be a deep neural network model, loosely inspired by CDZs (Meyer and Damasio, 2009) described in §2.2. The deep neural network will encode the high dimensional sensorial input data into layers with gradually reduced dimensions, up to a final low dimensional feature space, reminding the convergent process in Meyer & Damasio’s CDZs. This small feature space, in turn, becomes the input of a similar neural network acting as decoding, in which the dimensions of the layers are gradually increasing, paralleling the divergence in CDZs. Note that, when the desired output of the decoding network is the reconstruction of the original input, the combination is an *autoencoder*, Hinton’s idea that gave birth to deep learning itself, before being called “deep learning” (Hinton and Salakhutdinov, 2006; Hinton et al., 2006). Autoencoder has then become one of the most effective way of generating new data with deep learning (Dosovitskiy et al., 2017). In our case there will be both a decoding network reconstructing perceptual data, which plays the role of imagery, and a different network which decodes in the format of motor controls.

3.2 Action Selection

This component corresponds roughly to the processes performed in the brain by basal ganglia, typically in coordination with DLPFC, as described in §2.3. It has the function of deciding the agent’s behavior, and can accept high-level biases, such as automation level or driving style. Its implementation will be based on the multihypothesis sequential probability ratio test, an asymptotically optimal statistical test for decision making that has been shown to be a possible computational abstraction of the decision function performed by the cortex and basal ganglia.

3.3 Forward Model

This last component performs a role that is supposed to take place in the cerebellum, as discussed in §2.3: forward model. The input of this component is made up of efferent copies of motor commands, and the output is in term of intermediate perceptual data of the scene, as realized if the motor commands have been issued. This way, it would be possible to detect salient, and potentially dangerous situations, without having actually experience those situations in reality.

4 SYSTEM ARCHITECTURE

The system here proposed implements three different environments: the "wake" state, the "dream" state and quality assurance. Fig 2 shows an outline of this architecture.

4.1 The Wake State

This state corresponds with real driving, during which the agent records situations that are considered worthy to be re-enacted in the dream state. Those salient situations are determined according to several criteria. The most relevant events occur when there is a discrepancy between the predictions of the internal model and what actually happens in the real world. These situations can point out some imperfections in the prediction/planning model, or are the indication of the occurrence of a novel condition the agent has no knowledge about. Other relevant criteria that are taken into account are space-time separation with other vehicles, jerkiness of control, compliance with traffic rules, and also traffic and energy efficiency.

In order to be able to re-run experienced situations, the agent notes any event considered worthy of further analysis. The agent logs both the low-level sensory and control signals, and the high-level signals represented by the internal states of the agent's architecture. In such way, it is possible to record the "intentions" of the agent (or the estimated intentions of the human driver) at all levels of the sensorimotor system. Hence, the agent is able to run simulations of alternative lower-level strategies while preserving the higher-level intentions.

4.2 The Dream State

During this state, the agent can explore a simulated world to learning new behaviors, by testing different situations. The state is implemented in the open-source open-source virtual reality driving simulation environment OpenDS (www.opends.eu), based primarily on multibody systems (MBS) technology (Blundell and Harty, 2004), with some machine learning sub-model extensions. MBS is a general way for engineers to simulate large-scale physical system dynamics, including complex systems of bodies under the action of external forces, control loops, and other conditions. One significant advantage of MBS, compared to e.g., cerebellar forward models, is longer-term accuracy deriving from the universal physical principles substrate.

MBS can be instantiated with a very large number of parameters representing physical quantities. Selecting the model parameters from a distribution can be a method for generating imaginary conditions. Nevertheless, the distribution values must be plausible (it does not make sense to simulate a road with a friction coefficient equal to 10 since this never happens in reality). Thus, the distributions can either derive from the agent's observations whilst in the "wake" state or from a-priori knowledge, when it is is possible without loss of generality.

A radical different approach for generating imaginary situations is to use the same deep neural network model at the core of the inverse model of the agent, described in §3.1. This model has one principal decoding part that generates detailed motor commands from features at low dimensions, but also a secondary decoding part which works exactly as an autoencoder. It is therefore possible to generate a number of imaginary situations by randomly exploring this small feature space, with the caution of constraining these random combinations towards those situations that are more likely to be useful.

Another possible method is via generalized, top-down "motor babbling", a well established notion in the framework of perception-action machine learning (Shevchenko et al., 2009; Windridge et al., 2012; Windridge, 2017) and applied in cognitive robotics (Dearden and Demiris, 2005). It is akin to the process of language learning in infant humans: starting from minimal bootstrap sets of percepts and actions, the entire action space is randomly sampled. For each new motor action that produces a discernible perceptual output in the current perceptual set, the produced percept is allocated. This process can be carried out at various levels in the perception-action hierarchy.

4.2.1 Optimization

The final purpose of the simulation environment is the development of optimized behaviors (at all levels). One of the possible strategy is *optimal control* (OC), a methodology for producing optimal solutions

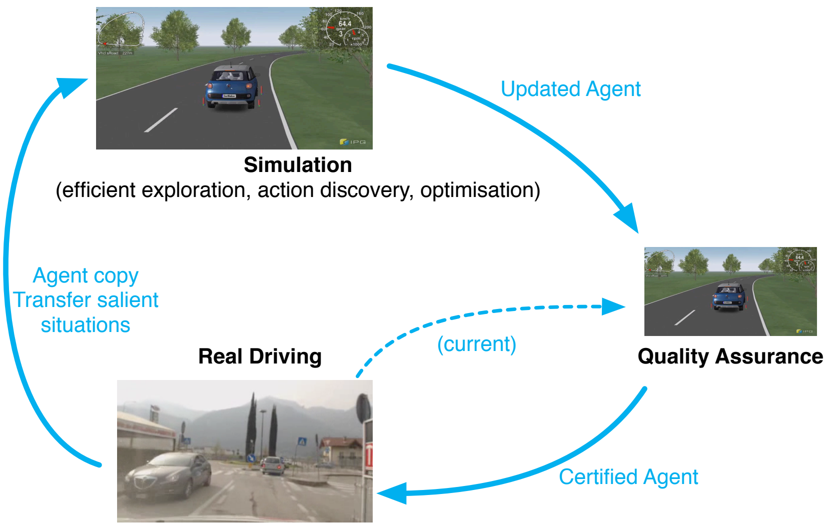

Figure 2: The system architecture, consisting of three components: real driving, simulation/optimization and quality assurance.

for goal-directed problems (Bertolazzi et al., 2005). Despite being OC a methodological framework based on differential algebraic equations and other mathematical structures, there are empirical evidences that the planning of hand movement to reach a target, as performed by the brain, corresponds to the minimization of the integral of the squared jerk, as in OC (Liu and Todorov, 2007). Therefore, OC too is in line with the brain-inspired design principle of this project.

Another possible approach is *motivated learning*, inspired by the learning mechanisms in animals, and thus better suited for higher-level strategies. In this case, the agent takes advantage of opportunistic interactions with the environment to develop a knowledge of what actions cause predictable effects in the environment. In this way, the agent is able to build internal models of action-outcome pairings in the brain.

4.3 Quality Assurance

Lastly, the third environment shown in Fig. 2 has the purpose of quality assurance. As the previous "dream" state, this environment is based on a multibody simulation system too, this one called CarMaker by IPG Automotive (ipg-automotive.com).

By testing against several test cases, the system certifies new versions of the agent, ensuring that any updated agent works no worse than the previously optimal one. It also helps to identify over-fitting, such as when the agent learns to better cope with the most recently-dreamed situation at the expenses of earlier ones.

5 CONCLUSION

This position paper has introduced a biologically inspired layered control architecture for intelligent vehicles, which relies on the recent developments in cognitive systems. We believe that the adoption of hierarchical perception-action learning via a dream simulation mechanism can considerably extend the utility of training data during learning. Moreover, this approach seems to offer advantages in terms of resources needed for development.

ACKNOWLEDGEMENTS

This work is supported by the European Commission under Grant 731593 (Dreams4Cars).

REFERENCES

Bertolazzi, E., Biral, F., and Da Lio, M. (2005). Symbolic-numeric indirect method for solving optimal control problems for large multibody systems: The time-optimal racing vehicle example. *Multibody System Dynamics*, 13(2):233–252.

Blundell, M. and Harty, D. (2004). *The multibody systems approach to vehicle dynamics*. Elsevier, Amsterdam.

Bogacz, R. and Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, 19(2):442–477.

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1585–1599.

Cisek, P. and Kalaska, J. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, 33:269–298.

Da Lio, M., Thill, S., Svensson, H., Gurney, K., Anderson, S., Windridge, D., Yksel, M., Saroldi, A., Andreone, L., and Heich, H.-J. (2017). Exploiting Dream-Like Simulation Mechanisms to Develop Safer Agents for Automated Driving. Yokohama, Japan.

Dearden, A. and Demiris, Y. (2005). Learning forward models for robots. *International Joint Conference on Artificial Intelligence*, 19:2–7.

Dixit, V. V., Chand, S., and Nair, D. J. (2016). Autonomous Vehicles: Disengagements, Accidents and Reaction Times. *PLOS ONE*, 11(12):e0168054.

Dosovitskiy, A., Springenberg, J. T., Tatarchenko, M., and Brox, T. (2017). Learning to generate chairs, tables and cars with convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):692–705.

Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks*, 12(78):961–974.

Gallagher, S. (2007). Social cognition and social robots. *Pragmatics & Cognition*, 15(3):435–453.

Gibson, J. J. (1979). *The Ecological Approach to Perception*. Houghton Miflin, Boston (MA).

Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. *The Behavioral and brain sciences*, 27(3):377–96; discussion 396–442.

Hesslow, G. (2012). The current status of the simulation theory of cognition. *Brain research*, 1428:71–9.

Hinton, G. E., Osindero, S., and Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554.

Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 28:504–507.

Jeannerod, M. and Frak, V. (1999). Mental imaging of motor activity in humans. *Current Opinion in Neurobiology*, 9(6):735–739.

Kalra, N. and Paddock, S. M. (2016). Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice*, 94:182–193.

Lewis, J., Chambers, J. M., Redgrave, P., and Gurney, K. N. (2011). A computational model of interconnected basal ganglia-thalamocortical loops for goal directed action sequences. In *BMC Neuroscience*, volume 12(Suppl 1), page 136, Stockholm.

Liu, D. and Todorov, E. (2007). Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27(35):9354–68.

Meyer, K. and Damasio, A. (2009). Convergence and divergence in a neural architecture for recognition and memory. *Trends in Neurosciences*, 32(7):376–382.

Moulton, S. T. and Kosslyn, S. M. (2009). Imagining predictions: mental imagery as mental emulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521):1273–1280.

Redgrave, P., Prescott, T., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89:1009–1023.

Shevchenko, M., Windridge, D., and Kittler, J. (2009). A linear-complexity reparameterisation strategy for the hierarchical bootstrapping of capabilities within perceptionaction architectures. *Image and Vision Computing*, 27(11):1702–1714.

Thill, S. and Svensson, H. (2011). The inception of simulation: a hypothesis for the role of dreams in young children. In Carlson, L., Hoelscher, C., and Shipley, T. F., editors, *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, pages 231–236, Austin, TX. Cognitive Science Society.

Ungerleider, L. and Mishkin, M. (1982). Two cortical visual systems. In Ingle, D. J., Goodale, M. A., and Mansfield, R. J. W., editors, *Analysis of visual behavior*, pages 549–586. MIT Press, Cambridge (MA).

Windridge, D. (2017). Emergent Intentionality in Perception-Action Subsumption Hierarchies. *Frontiers in Robotics and AI*, 4.

Windridge, D., Felsberg, M., and Shaukat, A. (2012). A Framework for Hierarchical Perception-Action Learning Utilizing Fuzzy Reasoning. *IEEE transactions on systems man and cybernetics Part B Cybernetics a publication of the IEEE Systems Man and Cybernetics Society*, 43(1):155–169.

Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature reviews. Neuroscience*, 12(12):739–51.

Wolpert, D. M. and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7-8):1317–1329.

Wolpert, D. M., Miall, R. C., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in cognitive sciences*, 2(9):338–47.