Prediction Error-Driven Memory Consolidation for Continual Learning: On the Case of Adaptive Greenhouse Models

Guido Schillaci¹ · Uwe Schmidt² · Luis Miranda³

Received: 15 May 2020 / Accepted: 21 December 2020 / Published online: 27 January 2021
© The Author(s) 2021

Abstract

This work presents an adaptive architecture that performs online learning and faces catastrophic forgetting issues by means of an episodic memory system and of prediction-error driven memory consolidation. In line with evidence from brain sciences, memories are retained depending on their congruence with the prior knowledge stored in the system. In this work, congruence is estimated in terms of prediction error resulting from a deep neural model. The proposed AI system is transferred onto an innovative application in the horticulture industry: the learning and transfer of greenhouse models. This work presents models trained on data recorded from research facilities and transferred to a production greenhouse.

Keywords Adaptive models · Deep neural networks · Episodic memory · Memory consolidation · Greenhouse

GS has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 838861 (Predictive Robots). Predictive Robots is an associated project of the Priority Programme "The Active Self" of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). GS has also received funding from the EU-H2020 Open Cloud Research Environment (OCRE) project and the Marie Curie Alumni Association as part of the European Open Science Cloud initiative, for running this project on GPU accelerated computing resources on the cloud. GS and LM have received funding from the DFG project on initiation of international collaboration "Adaptive Architectures for Transferability of Greenhouse Models". LM received funding from the Project PROSIBOR. The project "Development of a sensor based intelligent greenhouse management system (PROSIBOR (FKZ: 2815701315))" is supported by funds of the Federal Ministry of Food and Agriculture (BMEL) based on a decision of the Parliament of the Federal Republic of Germany via the Federal Office for Agriculture and Food (BLE) under the innovation support programme. *Author contribution:* GS and LM designed and implemented the computational models, the experiments and part of the data pre-processing. LM and US collected and pre-processed the greenhouse data.

✉ Guido Schillaci
guido.schillaci@santannapisa.it

¹ The BioRobotics Institute and Department of Excellence in Robotics & AI, Scuola Superiore Sant'Anna, Pisa, Italy

² Department of Biosystems Engineering, Humboldt-Universität zu Berlin, Berlin, Germany

³ Faculty of Forest and Environment, Eberswalde University for Sustainable Development, Eberswalde, Germany

1 Introduction

Adaptivity is about adjusting behaviours or beliefs to achieve novel objectives or to respond to unexpected circumstances. Of crucial importance for biological systems, adaptivity is one of the most challenging capabilities to implement in artificial systems. Developmental robotics addresses this challenge by taking inspiration from models of human development and from principles of brain functioning. [2, 18]. Indeed, infant brains are continuously exposed to rich and novel sensorimotor experience while morphological and environmental conditions are changing. Skills acquired at a certain point in time—e.g. sitting up, manipulating toys—need to be re-adapted as the proportions of growing body parts change and as other capabilities emerge.

Brain science communities converge on considering the somatosensory cortex of the human brain as playing a role in the implementation of adaptive body representations [15]. These representations are formed along the rich sensorimotor information the individual is exposed to, while interacting with its surroundings. Research suggests that experienced sensorimotor contingencies and action-effect regularities are stored in the brain, allowing later processes of anticipation of sensorimotor activity. This has been shown to be crucial for adaptive behaviours, perception [12], motor control [1], memory [8, 11] and many other cognitive functions [14, 29], and has inspired a wide range of computational models for artificial systems [4, 7, 30]. However,

despite the promising results in robotics and AI, a number of challenges still remain open. Among these is the question about how adaptivity can be leveraged in *lifelong learning* systems. Although there is an increasing understanding of how biological systems balance the integration of new knowledge while retaining past experience in memory, an implementation of such strategies in artificial systems is still arduous.

In mammals, memory is composed of multiple systems supported by different structures in the brain [36]. One of these systems, i.e. episodic memory, is crucial for adaptive behaviours, as well as for other cognitive functions such as planning, decision-making and imagination [25]. Memory traces are stabilised in the brain after their initial acquisition through memory consolidation [37]. Consolidation occurs at different levels in the brain, including a faster, synaptic (hippocampal) level and a slower, more stable (neocortical) system level. System consolidation seems to be driven by the hippocampus, which reorganises its stored temporal and labile memories into more stable traces in the neocortex [36]. The rate of consolidation can also be influenced by the congruence between prior knowledge and the information that is going to be stored [38]. Recent studies suggest that if the information to be learned is consistent with prior knowledge, neocortical consolidation can be more rapid [19, 36]. In other words, the way memory is updated appears to be dependent on the extent new information is likely to be formed [9, 32, 33]. Consolidated memories are not static imprints of past experiences, but are rather malleable and can be updated or reconsolidated [16, 27, 34]. A key component of this process seems to be the capability of the brain to evaluate a prediction error, or a surprise signal, which would be necessary for destabilising and reconsolidating memories. Evidences suggest also that formation and consolidation of long-term memories occur during sleep, where experienced events are likely to be reactivated [3]. The rate of memory consolidation also depends on the developmental stage of the individual. Infants show weaker retention of experience compared to adults, thus reflecting a tendency of young brains to save more newly acquired experience [13].

The present work brings a twofold contribution to this special issue. Firstly, it advances the state-of-the-art on continual learning in artificial systems. In particular, it proposes an online learning framework implementing an episodic memory system, in which memories are retained according to their congruence with the prior knowledge stored in the system. Congruence is estimated in terms of prediction error resulting from a generative model.

Secondly, it shows that some of the paradigms of developmental robotics and of brain-inspired computational modelling can be transferred from laboratories to innovative applications. In particular, we apply this research in practical

horticulture: the design of greenhouse models for monitoring physiological parameters of plants—with the goal of increasing crop yield—and their transfer from research to production greenhouse facilities.

1.1 AI Transfer: Adaptive Greenhouse Models

Continual learning, i.e., the capability of a learning system to continually acquire, refine and transfer knowledge and skills throughout its lifespan, has represented a long standing challenge in machine learning and neural network research [26]. Training neural networks in an online and prolonged fashion without caution typically rises *catastrophic forgetting* issues [20]. Catastrophic forgetting consists of the overwriting of previously learned knowledge that occurs when a model is being updated with new information. Researchers have been trying to tackle this issue through different strategies [5, 17, 31]. These include consolidating past knowledge already present in a short-term memory system into a long-term memory one [21], or employing an episodic memory system [20] that maintains a subset of previously experienced training samples and replays them, along with the new samples, to the networks during the training. This paper adopts a mixed approach which uses episodic memory replay and prediction-error driven consolidation to implement online learning in deep recurrent neural networks. Importantly, this work aims at transferring this AI strategy onto an application for the innovative greenhouses industry.

Greenhouses are complex systems comprising technical and biological elements. Similarly to robots, their state can be measured and modified through control actions, for instance on the internal climate. Modelling the mappings between different sensors and control actions as well as the resulting measurements allows to anticipate the effects of an intervention upon the greenhouse conditions, to better plan further control actions and, ultimately, to increase crop yield. Several studies can be found in the horticulture literature showing that neural networks can model different processes occurring in a greenhouses, including internal climate [10] and yield [6, 28]. Experiments by [22] used multilayer perceptrons to predict time series in greenhouses, particularly leaf tissue temperature, transpiration and photosynthesis rates of a tomato canopy. The authors used chained simulations to generate predictions of several time steps, using 3 time steps for all input signals. A more thorough investigation of the time steps needed to predict time series inside a greenhouse is given by [24], who points that a static selection of time steps gives poor results after three historical steps (15 min) in the inputs. This is due to different time constants involved in the system, as several inputs show very fast variations while others appear

delayed. These experiments suggest that more elaborated models are needed to account for the different memory length needed to make a prediction. This length is different for each input and changes at times of the day and seasons of the year.

Despite their potential impact in several applications in the field, adaptive models have received little attention in the horticulture scientific community (e.g. [35]). Indeed, the possibility to adapt can facilitate the transfer of models from research facilities to the production greenhouses. In a preliminary study [23], we showed that a learning architecture characterised by deep recurrent neural networks and an episodic memory system can enable the portability of greenhouse models. The model exposed to a big amount of data recorded from a research greenhouse can be transferred to a production facility, requiring less training data from the new greenhouse setup. This approach can have a high impact on the greenhouse industry, as it would allow to design and train optimal models at research greenhouses and quickly re-adapt them to different production facilities and crops.

Here, we extend our previous study [23] by introducing a more efficient memory consolidation strategy and by providing a more comprehensive analysis of the different aspects of the architecture. As in the previous work, we train a computational model for estimating the transpiration and photosynthesis of a hydroponic tomato crop by using measurements of the climate. The models are trained and tested using data from two greenhouses in Berlin, Germany. Thereafter, the adaptive model is fed with data from a production greenhouse in southern Germany, near Stuttgart, where other tomato varieties were grown under different irrigation and climate strategies.

2 Methodology

The computational model adopted here consists of a deep neural network, in part composed by Long Short-Term Memory (LSTM) layers, characterised by two outputs—transpiration and photosynthesis—and a time series of six sensor values as input. In particular, climate data (air temperature, relative humidity, solar radiation, CO₂ concentration) and temperature of two leaves are used as sensor data. The model is used to predict transpiration and photosynthesis rates from the sequence of sensor data. Anticipating these information allows better control of the climate and, consequently, an increase of the yield. This aspect is however not covered by this study.

The samples have been pre-recorded from three different greenhouses (hereon: GH1, GH2 and GH3), with a rate of one multi-sensors measurement every 5 minutes. GH1 and GH2 are research greenhouses located in Berlin. Recordings

have been carried out during several years: 2011 to 2014 for GH1, and 2015 to 2016 for GH2. GH3 is a production greenhouse located near Stuttgart, Germany. Data from 2018 was obtained for this greenhouse.

We test two models¹, both with inputs consisting of *fixed-length time series* of six sensors data. The first model (M1) takes as input a window of 288 subsequent samples from the six sensors, corresponding to one full day of recordings, given that samples are captured every 5 minutes. The second model (M2) takes as input a window of 576 subsequent 6D samples, corresponding to two full days of recordings². The output consists of a 2D vector representing the transpiration and photosynthesis rates recorded at the final time step of the window.

Datasets are prepared so that input-output training samples can be sequentially extracted, to simulate an online learning process. For both models, the first training phase includes 5 cultivation years (2011 to 2014) from GH1. Subsequent phases use the cultivation years 2015 (GH2), 2016 (GH2) and 2018, regarding the commercial greenhouse (GH3). In all cases, the time series are truncated during the winter production pauses.

For model M1, this results in 26197 training samples from GH1 exposed sequentially to the learning process. After all samples are covered, the model is exposed to 7079 samples from GH2, (2015) and to 5566 samples from GH2 (2016). Finally, the model is exposed to 1153 samples from GH3 (2018). During each of these training phases, performance of the learning system is estimated by computing the mean squared error (MSE) on test datasets extracted from the corresponding greenhouse. In particular, test datasets consist of 1377 samples (1/20th of the GH1 training dataset size) for GH1, 372 samples for GH2 (2015), 292 samples for GH2 (2016) and finally 60 samples for GH3.

In another experiment, model M2 is trained and tested on smaller datasets, defined by wider input windows (two days, or 576 samples). In this experiment, M2 is exposed, in sequence, to 24949 training samples (tested on 1311 samples) from GH1, to 6831 training samples (tested on 359 samples) from GH2 (2015) and to 5431 training samples (tested on 285 samples) from GH2 (2016), and finally to 1096 training samples (tested on 57 samples) from GH3 (2018).

¹ Diagrams of the structure of the models, as well as the source code, can be found here: https://github.com/guidoschillaci/online_lstm_episodic_memory.

² One day included a full daylight-night shift in the data. We avoided dealing with smaller portion of time windows, and chose instead multiples of one day. Considering the good performance of the system using two days input data, as described in the results section, we did not increase anymore the input size to allow a clearer differentiation between memory consolidation strategies.

Test data are not included in the training sets³.

Model updates are performed on batches of 32 subsequent samples. As discussed above, an episodic memory system is used to reduce catastrophic forgetting issues. This system replays samples (together with the current batch) when updating the model's weights. Samples observed over time are stored into an episodic memory and retained following a prediction-error driven consolidation scheme: a mechanism that chooses which samples to maintain in the episodic memory based on their expected contribution to the learning progress. Each memory element consists of an input-output mapping, i.e. a fixed-length time series (of one day for model M1; of two days for model M2) of 6D vectors as input and a 2D vector as output. A memory element is also characterised by a prediction error—i.e. how the model's guess about such a stored experience deviates from the actual measured value – and by an expected learning progress—estimated as the absolute value of the derivative of two subsequent prediction errors—. More precisely, the learning progress LP is calculated as:

$$LP = |\epsilon_t - \epsilon_{t-1}|$$

$$= |(s_t^* - s_t) - (s_{t-1}^* - s_{t-1})|$$

where ϵ is the prediction error calculated as the Euclidean distance between the sensory state s (transpiration and photosynthesis) and the sensory prediction s^* . Sensory predictions are inferred by feeding the 6D input of a memory element into the model. After each model update, the derivative of the prediction error associated to each memory element is updated.

For both models M1 and M2, we compare three different memory consolidation strategies and one that does not adopt any episodic memory system—hereon named *no-memory* strategy. The first strategy, i.e., *discard high LP*, tends to consolidate memory elements that produced little variations in the prediction error. This is performed by discarding, at every memory update, the element characterised by the highest absolute value of the derivative of the prediction error (an estimate of the expected contribution to the learning progress) and by replacing it with the most recently observed sample. A second strategy, hereon named *discard low LP*, tends to consolidate memory elements that produced big variations in the prediction error, likely to impact more on the learning progress during the next training iteration. In particular, it discards the memory element characterised by the smallest variation in the prediction error. This strategy is more in line with the literature reviewed at the beginning of

Table 1 The configurations of the experiments carried out in this work. Each experiment is run 10 times

Design of experiments			
ID	Model	Mem. consolidation	Update prob.
1	M1	No memory	–
2	M1	Discard high LP	0.05%
3	M1	Discard low LP	0.05%
4	M1	Discard random	0.05%
5	M1	No memory	–
6	M1	Discard high LP	0.4%
7	M1	Discard low LP	0.4%
8	M1	Discard random	0.4%
9	M2	No memory	–
10	M2	Discard high LP	0.05%
11	M2	Discard low LP	0.05%
12	M2	Discard random	0.05%
13	M2	No memory	–
14	M2	Discard high LP	0.4%
15	M2	Discard low LP	0.4%
16	M2	Discard random	0.4%

this paper, and we expect it to outperform the others. A third baseline strategy, named *discard random*, implements the standard memory consolidation approach in machine learning: at every memory update a randomly chosen sample is discarded from the memory.

Fig. 1 Mean squared error over time of the following experiments: ► model M1 (input: one day of observations) in the *stable* memory update configuration (first column); model M1 in the *plastic* memory update configuration (second column); model M2 (input: two days of observations) in the *stable* memory update configuration (third column); model M2 in the *plastic* memory update configuration (fourth column). The first row shows the MSE when no episodic memory is employed; the second and third rows show the MSE of the models employing the *discard high LP* and *discard low LP* memory consolidation strategies, respectively; the fourth row shows the MSE of the model employing the baseline *discard random* consolidation strategy. Vertical axes indicate MSE values in the logarithmic scale. Horizontal axes represent time, in form of the iterations at which MSE has been estimated. Model update is performed every time a 32-batch of samples is observed. MSE is not computed at every model update, but rather at a slower pace, i.e. every four model updates. Vertical dashed lines indicate switches between training datasets. From time 0 to the iteration marked with the red vertical dashed line, the model is exposed to data recorded from GH1. From the iteration marked with the red line to the purple one, the model is exposed to data from GH2 (2015). From the iteration marked with the purple line to the blue one, the model is exposed to data recorded from GH2 (2016). Finally, from the instant marked with the blue line until the end, the model is exposed to data recorded from the production greenhouse GH3 (2018). Each experiment is repeated 10 times. Solid lines show the average MSE over the 10 runs calculated on four different test datasets (green plot: data from GH1, red plot: data from GH2, year 2015; purple plot: data from GH2, year 2016; blue plot: data from GH3). Shaded areas indicate errors (mean $\pm$ std.dev.) over the 10 runs

³ The number of training samples available for models M1 and M2 vary as a result of selecting windows of one or two days, the step size used to scan the original dataset and the random extraction of test data. The extraction algorithm can be found in the *loaddataset.py* script in the github repository.

Finally, we compare the different architectures varying another hyper-parameter, i.e., the probability of updating the memory: in a *stable* configuration, the memory is updated 5% of the times a new sample is observed; in a *plastic* configuration, this probability is set to 40%, therefore updating the memory much more frequently than in the *stable* setup. Table 1 summarises the experiments.

3 Results

Fig. 1 shows the mean squared error over time for each of the experiments depicted in Table 1.

The absence of an episodic memory system (experiments 1, 5, 9, 13) produces higher MSE values and big fluctuations in the MSE curves, likely due to catastrophic forgetting issues. Abrupt deterioration of system performance can be

Table 2 Quantitative analysis comparing the *discard low LP* and *discard random* consolidation strategies

	P1	P2	P3	P4
T1	-	-	-	*
T2	-	-	**	-
T3	-	-	-	*
T4	-	-	-	-

Exp.#3 vs Exp.#4

	P1	P2	P3	P4
T1	**	**	**	**
T2	-	**	*	-
T3	-	-	-	**
T4	-	-	-	-

Exp.#7 vs Exp.#8

	P1	P2	P3	P4
T1	-	*	-	-
T2	-	-	*	-
T3	-	-	**	-
T4	-	-	-	-

Exp.#11 vs Exp.#12

	P1	P2	P3	P4
T1	-	-	**	-
T2	-	-	**	-
T3	-	-	-	-
T4	-	-	-	-

Exp.#15 vs Exp.#16

Periods P1, P2, P3, P4 indicate the learning phases characterised by the different training data: GH1 (2011 to 2014), GH2 (2015), GH2 (2016) and GH3 (2018), respectively. Rows indicate the test dataset used to calculate the MSE. A linear regression is calculated on each MSE curve segment, i.e., on the MSE values within the indicated training period. A green cell indicates that the MSE slope of the (*discard low LP*) strategy is smaller than that of (*discard random*). That is, a steeper descent in the MSE is observed using the *discard low LP* than the *discard random* one. Inversely, red cells indicate that *discard random* strategy produced a steeper MSE descent, compared to the *discard low LP* strategy. Gray cells refer to test data from a GH not yet used as training source, and are hence omitted. Asterisks indicate p-values of the interaction analysis, and thus of the slope difference: one asterisk indicates p-value ≤ 0.05 ; two asterisks indicate p-value ≤ 0.01 . Empty values (-) indicate that the difference was not statistically significant

observed whenever training datasets are switched (see the peaks in the MSE near the vertical dashed lines), showing the poor adaptive capabilities of the model⁴.

By contrast, an episodic memory system produces a more stable learning progress (see Fig. 1). Overall, the *discard low LP* memory consolidation strategy outperforms the other methods, as expected. A model under this configuration presents more stability after the changes in the training distributions. Table 2 presents a quantitative analysis in support to these statements. In particular, we analysed the difference between the slopes of the linear regression computed on the MSE produced by the *discard low LP* and *discard random* consolidation strategies⁵. The statistical significance of the slope differences is estimated by means of an interaction analysis⁶. Overall, the *discard low LP* has a tendency to over-perform the *discard random* strategy. *Discard random* brought better results—marked in red in Table 2—in fewer cases than for the *discard low LP*. They occur in plastic models (40% memory update probability) and only in GH1: the largest and first training group in all experiments⁷.

The *discard high LP* strategy over-consolidates past and, perhaps, less informative experiences (see later comment about the variance of the stored episodic memories). This can be noticed in Fig. 2, which illustrates the content of the episodic memory over time for the *stable* (left column) and *plastic* (right column) configurations. The plots show the amount of elements from GH1 (green), GH2 (red) and GH3 (blue) stored in the memory at each time⁸. Notably, the *discard low LP* strategy fills up the memory with new samples faster than *discard high LP*. A similar trend is observable in the *discard random* strategy. Nonetheless, the *discard low LP* strategy maintains memory elements from previous greenhouses longer than the *discard random* strategy, likely affecting models' performance (1). For example, in experiments 15 and 16, MSE on the GH2 (2015) test dataset, during exposition of the models to GH2 (2016) training data, decreases faster in the *discard random* than in the *discard low LP* strategy. A similar situation can be observed in Exp. 7 and 8.

Replaying more recent samples during the model update is likely to increase the plasticity of the system. In fact, smaller peaks in the MSE in the *discard low LP* plots can be observed when the distribution changes. *Plastic* configurations in general respond faster to new data (see final training instances—blue curves—in Exp. 7, 8, 15 and 16). Moreover, the *discard low LP* strategy ensures that a higher variance in the values stored in the memory is maintained over time,

⁴ Spikes of the MSE in the *no-memory* strategy already occur before a switch to a new dataset is performed. We believe this to be a seasonal effect in the data, as the time series are truncated during winter production pauses (see Sect. 2).

⁵ We omit the other, less interesting, strategies, as trends can be qualitatively inferred from Fig. 1.

⁶ For each period following a dataset change, a linear model is trained and p-values of the interaction between time and group are reported in Table 2.

⁷ We believe that analysing the impact of the number of elements to be consolidated each time or of the moment when to integrate new observations, may help in explaining these trends.

⁸ The maximum size of the memory is set to 500 elements in all the experiments. The memory is filled up with any observed sample, until it is full. Thereafter, the chosen consolidation strategy is applied. We tested different memory sizes (in the range 100–1000, empirically chosen). We decided to fix the memory size to 500 as it allowed easier visualisation of the system performances.

Fig. 2 Content of the episodic memory over time for memory consolidation strategies (means $\pm$ std.dev. over 10 runs per experiment) using model M1. y-values represent the amount of elements (from 0 to 500, where 500 is memory size) from each dataset in the memory.

Horizontal axes represent time. The left column shows *stable* configurations, while the right column shows *plastic* configurations. The plots for the model M2 have been omitted, as they closely resemble those of M1

Fig. 3 The variance of the content of the memory over time

compared to the other strategies (see Fig. 3). We believe that this helps the model to maintain a good balance between stability and plasticity.

Finally, it is worth noting that a principal component analysis (PCA) carried out on all the datasets (Fig. 4), and estimated on 8 dimensions—i.e., six sensors data, and transpiration and photosynthesis—shows a partial overlap between datasets. GH2 and GH3 data seem to be partly represented by GH1 data. As can be seen in Fig. 1, the MSE curves for GH3 test data (blue curves) during the first

learning phase—i.e., where data from GH1 is used—show a steeper descent compared to the others, although no data from GH3 is yet being learned by the model. This can be explained by the fact that GH1 was used to test a number of climate control strategies, resulting in a broader range of conditions being reflected in the data. Additionally, greenhouses GH1 and GH2 share the same construction and location, while GH3 is bigger and subject to different meteorological conditions. Despite similarities in the datasets, we believe that the proposed memory consolidation strategies

Fig. 4 Principal component analysis of all datasets

based on prediction error estimates can be used to produce more stable learning systems, compared to standard consolidation strategies.

4 Conclusions

This paper presented an architecture in which episodic memory replay and prediction-error driven consolidation are used to tackle online learning in deep recurrent neural networks. Inspired by evidences in brain sciences, memories are retained depending on their congruence with the prior knowledge stored in the system. Congruence is estimated in terms of prediction error resulting from a generative model, a deep recurrent neural network. This approach produces a good balance between stability and plasticity in the model and tends to outperform standard memory consolidation strategies.

Importantly, this work aimed at transferring developmental robotics solutions onto an application for the greenhouse industry, i.e., the transfer of climate models from research facilities to production greenhouses. We show that the system exposed to data recorded from a research greenhouse can be transferred to a production facility, without facing the need to re-train on a big amount of data from the new setup, a process that is costly and involves a high risk of damaging the crop.

Funding Open Access funding provided by Scuola Superiore Sant'Anna

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adams RA, Shipp S, Friston KJ (2013) Predictions not commands: active inference in the motor system. *Brain Struct Func* 218(3):611–643
- Asada M, Hosoda K, Kuniyoshi Y, Ishiguro H, Inui T, Yoshikawa Y, Ogino M, Yoshida C (2009) Cognitive developmental robotics: a survey. *IEEE TAMD*
- Born J, Wilhelm I (2012) System consolidation of memory during sleep. *Psych. Res* 76(2):192–203
- Cangelosi, A., Schlesinger, M.: *Developmental robotics: From babies to robots*. MIT press (2015)
- Chen, Z., Liu, B.: *Lifelong machine learning*. Synthesis lectures on AI and machine learning (2018)
- Ehret DL, Hill BD, Helmer T, Edwards DR (2011) Neural network modeling of greenhouse tomato yield, growth and water use from automated crop monitoring data. *Comp. Electr. Agriculture* 79(1):82–89
- Eppe, M., Magg, S., Wermter, S.: Curriculum goal masking for continuous deep reinforcement learning. In: *IEEE ICDL-EpiRob*, pp. 183–188 (2019)
- Ergo K, De Loof E, Verguts T (2020) Reward prediction error and declarative memory. *Trends in Cog. Sci*
- Exton-McGuinness MT, Lee JL, Reichelt AC (2015) Updating memories-the role of prediction errors in memory reconsolidation. *Behav Brain Res* 278:375–384
- Fitz-Rodríguez E, Kacira M, Villarreal-Guerrero F, Giacomelli GA, Linker R, Kubota C, Arbel A (2012) Neural network predictive control in a naturally ventilated and fog cooled greenhouse. *Acta Horticulturae* 952:45–52
- Fountas Z, Sylaidi A, Nikiforou K, Seth AK, Shanahan M, Roseboom, W (2020) A predictive processing model of episodic memory and time perception. *bioRxiv* (2020)
- Friston K (2012) Prediction, perception and agency. *Int J Psychophysiol* 83(2):248–252
- Gómez, R.L.: Do infants retain the statistics of a statistical learning experience? insights from a developmental cognitive neuroscience perspective. *Phil Trans Royal Soc B* 372(1711) (2017)
- Hohwy J (2013) *The predictive mind*. Oxford Un. Press,
- Holmes NP, Spence C (2004) The body schema and multisensory representation (s) of peripersonal space. *Cognitive Process* 5(2):94–105
- Jang AI, Nassar MR, Dillon DG, Frank MJ (2019) Positive reward prediction errors during decision-making strengthen memory encoding. *Nat Hum Beh*
- Kirkpatrick J, Pascanu R, Rabinowitz N, Veness J, Desjardins G, Rusu AA, Milan K, Quan J, Ramalho T, Grabska-Barwinska A et al (2017) Overcoming catastrophic forgetting in neural networks. *Proc Natl Acad Sci* 114(13):3521–3526
- Krichmar, J.L.: A neurobiologically inspired plan towards cognitive machines. In: *Towards Conscious AI Sys.* (2019)

19. McClelland JL (2013) Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Exp Psy*
20. McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102(3):419
21. Mermillod M, Bugaiska A, Bonin P (2013) The stability-plasticity dilemma: investigating the continuum from catastrophic forgetting to age-limited learning effects. *Front Psychol* 4:504
22. Miranda L, Lara B, Rockschi T, Dannehl D, Schmidt U (2019) Long-term prediction of biosignals from greenhouse-grown tomato. *Acta Horticulturae* 1242:739–746
23. Miranda L, Schillaci G.: An adaptive architecture for portability of greenhouse models. *Acta Horticulturae* (In Press—Preprint available at: <https://arxiv.org/pdf/1908.01643>)
24. Miranda Trujillo, L.C.: Artificial neural networks in greenhouse modelling : two modelling applications in horticulture (2018)
25. Murty VP, FeldmanHall O, Hunter LE, Phelps EA, Davachi L (2016) Episodic memories predict adaptive value-based decision-making. *Exp sych*
26. Parisi GI, Kemker R, Part JL, Kanan C, Wermter S (2019) Continual lifelong learning with neural networks: a review. *Neural Netwk*
27. Roscow EL, Jones MW, Lepora NF Behavioural and computational evidence for memory consolidation biased by reward-prediction errors. *bioRxiv* (2019)
28. Salazar R, López I, Rojano A, Schmidt U, Dannehl D (2014) Tomato yield prediction in a semi-closed greenhouse. In: *IHC2014*, pp. 263–270
29. Schillaci G, Hafner VV, Lara B (2016) Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents. *Front Robot AI* 3:39
30. Schillaci, G., Pico, A., Hafner, V.V., Hanappe, P., Coliaux, D., Wintz, T.: Intrinsic motivation and episodic memories for robot exploration of high-dimensional sensory spaces. *Sage Adaptive Behaviour* (2020)
31. Shin H, Lee JK, Kim J (2017) Continual learning with deep generative replay. In: *NeurIPS*
32. Simon KC, Gómez RL, Nadel L, Scalf PE (2017) Brain correlates of memory reconsolidation: a role for the tpj. *Neurobiology of learning and memory*
33. Sinclair AH, Barense MD (2018) Surprise and destabilize: prediction error influences episodic memory reconsolidation. *Learn Memory* 25(8):369–381
34. Sinclair AH, Barense MD (2019) Prediction error and memory reactivation: How incomplete reminders drive reconsolidation. *Trends Neurosci*
35. Speetjens SL, Stigter JD, Van Straten G (2009) Towards an adaptive model for greenhouse control. *Comput Electron Agriculture* 67(1–2):1–8
36. Squire LR, Genzel L, Wixted JT, Morris RG (2015) Memory consolidation. *Cold Spring Harbor perspectives in biology* 7(8):a021766
37. Traversa FL, Pershin YV, Di Ventra M (2013) Memory models of adaptive behavior. *Neural Networks*
38. Van Kesteren MT, Ruiter DJ, Fernández G, Henson RN (2012) How schema and novelty augment memory formation. *Trends Neurosci* 35(4):211–219