

# Learning with Chemical versus Electrical Synapses Does it Make a Difference?

Mónika Farsang<sup>1\*</sup>, Mathias Lechner<sup>2</sup>, David Lung<sup>1</sup>, Ramin Hasani<sup>2</sup>, Daniela Rus<sup>2</sup>, Radu Grosu<sup>1</sup>

**Abstract**—Bio-inspired neural networks have the potential to advance our understanding of neural computation and improve the state-of-the-art of AI systems. Bio-electrical synapses directly transmit neural signals, by enabling fast current flow between neurons. In contrast, bio-chemical synapses transmit neural signals indirectly, through neurotransmitters. Prior work showed that interpretable dynamics for complex robotic control, can be achieved by using chemical synapses, within a sparse, bio-inspired architecture, called Neural Circuit Policies (NCPs). However, a comparison of these two synaptic models, within the same architecture, remains an unexplored area. In this work we aim to determine the impact of using chemical synapses compared to electrical synapses, in both sparse and all-to-all connected networks. We conduct experiments with autonomous lane-keeping through a photorealistic autonomous driving simulator to evaluate their performance under diverse conditions and in the presence of noise. The experiments highlight the substantial influence of the architectural and synaptic-model choices, respectively. Our results show that employing chemical synapses yields noticeable improvements compared to electrical synapses, and that NCPs lead to better results in both synaptic models.

## I. INTRODUCTION

Continuous-time recurrent neural networks (CT-RNNs) are well-suited for robotic and control applications, because they are able to learn dependencies and patterns occurring in CT processes [1], [2], [3], [4]. These networks are described by ordinary differential equations (ODEs) representing the time evolution of the activation of neurons over time, based on the inputs and the internal dynamics of the network.

Most robotic and control applications are safety-critical. Their faulty operation may cause human harm or physical damage. Interpretable machine learning methods are thus a natural fit for robotic applications, as they allow the system behavior to be studied and audited by humans.

Liquid-time-constant neural networks (LTCs), a class of biology-inspired neural models with chemical synapses [4], have recently demonstrated interpretable-control capabilities in autonomous-driving and drone-navigation tasks [5], [6], [7], when organized in a sparse 4-layer Neural Circuit Policy architecture (NCP) [8]. However, the individual contributions of the NCP architecture and the synaptic model to interpretable control, is still poorly understood.

In this paper, we first show that electrical biologic synapses underlay CT-RNNs [2], and that chemical synapses

underlay CT-RNNs with synapse-bound activations, and input-dependent forget gates [9]. We then study how much learning with electrical versus chemical synapses really makes a difference in the network's performance when they share the same type of architecture. For a systematic comparison, we experiment with LTCs [4] and classic CT-RNNs [2], in both NCP [5] and fully-connected architectures.

We evaluate the performance of these networks in the complex task of autonomous car-navigation across different seasons. We first train the networks using imitation learning [10], where they learn to imitate the actions of a human driver through expert demonstrations. Subsequently, we deploy the trained networks in closed-loop settings to examine their adaptivity to previously unseen scenarios and their ability to react to environmental feedback. These evaluations are conducted using the data-driven photorealistic autonomous driving simulator VISTA [11]. This has been proven to yield consistent results, when comparing trained agents in closed-loop evaluations in simulation, and real-world tests.

Our main contributions are the following:

- We show that saturated electrical synapses are the foundation of classic continuous-time RNNs, and that saturated chemical synapses underlay CT-RNNs with synapse-bound activations and input-dependent forget gates.
- We conduct a comparative analysis between electrical and chemical synaptic models, respectively, in sparse NCP and fully-connected wiring architectures, evaluating their performance, interpretability and robustness.
- Our results show that: (1) The architectural wiring choice significantly impacts the network performance, with NCPs exhibiting better results with both synaptic models, and (2) Employing chemical synapses yields noticeable improvements compared to electrical synapses.

This paper is structured as follows: In Section II we discuss the bio-inspired models of electrical and chemical synapses, and the wiring model of NCPs. Section III provides a summary of related work on imitation learning, lane-keeping, and applications of the models under investigation. In Section IV, we outline the metrics of interest and detail the experimental setup. The results are presented in Section V, followed by a concluding discussion in Section VI.

## II. SYNAPTIC MODELS

The dynamic behavior of neurons is captured in neuroscience with electrical equivalent circuits (EECs) [12], [13].

<sup>1</sup>CPS, Technische Universität Wien (TU Wien), Austria

<sup>2</sup>CSAIL, Massachusetts Institute of Technology (MIT), MA, USA

\*monika.farsang@tuwien.ac.at

\*\*This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101034277.

In EECs, a neuron- $i$ 's membrane is regarded as a capacitor whose potential  $x_i$  is determined by the difference between intra- and extra-cellular concentrations of ions, respectively. The membrane potential thus evolves as follows:

$$C \dot{x}_i = i_{li} + i_{si}, \quad i_{li} = g_{li}(e_{li} - x_i) \quad (1)$$

where  $C$  is the membrane's capacitance, and  $i_{li}$  and  $i_{si}$  are the passive-leaking and dynamic-synaptic currents, traversing the capacitor, respectively. Without loss of generality, we will assume that  $C = 1$  in the rest of the paper. The leaking current, is determined by the Ohm's equation, where  $e_{li}$  is the resting (or reversal) potential of the neuron, and  $g_{li}$  is its leaking conductance. They are considered to be constant.

1) *Electrical synapses*: occur in all organisms, especially in embryonic stage. Their prevalent form is that of an Ohmic current  $i_{s,ji} = g_{ji}(y_j - x_i)$ , where  $y_j$  is either the potential  $x_j$  of a presynaptic neuron, or the potential  $u_j$  of an input. Constant  $g_{ji}$  is the conductance of the synapse. Hence, in electrical synapses,  $x_i$  faithfully follows  $y_j$ . Let  $y = [x, u]$ ,  $|x| = m$  and  $|u| = n$ . As a network, we thus obtain:

$$\dot{x}_i = g_{li}(e_{li} - x_i) + \sum_{j=1}^{m+n} g_{ji}(y_j - x_i) \quad (2)$$

Rearranging this equation in terms of  $x_i$  and  $e_{li}$ , one obtains:

$$\dot{x}_i = -(g_{li} + \sum_{j=1}^{m+n} g_{ji}) x_i + (g_{li} + \sum_{j=1}^{m+n} \frac{g_{ji}}{e_{li}} y_j) e_{li} \quad (3)$$

Equations (2)-(3) are the EECs of neural networks with electrical synapses, only. However, neurons saturate their overall conductance, through diverse cellular mechanisms. A more faithful and numerically stable representation is thus:

$$\dot{x}_i = -S(g_{li} + \sum_{j=1}^m g_{ji}) x_i + T(g_{li} + \sum_{j=1}^m h_{ji} y_j) e_{li} \quad (4)$$

where  $S$  and  $T$  are sigmoidal and hyperbolic-tangent functions, respectively, and  $h_{ji} = g_{ji} / e_{li}$ . This equation is normalized, with conductances  $g_{li}, g_{ji} \in [0, 1]$ . Hence, saturation is with  $S$ . Since  $h_{ij} \in [-1, 1]$  for  $e_{li} \in [-1, 1]$ , we use  $T$ .

Taking  $\tau = S(g_{li} + \sum_{j=1}^m g_{ji})$  one obtains the classic CT-RNNs [2], with constant gate  $\tau$ . Thus, classic CT-RNNs are neural networks with electrical synapses, only!

2) *Chemical synapses*: occur in all organisms, especially in later stages. A current  $i_{s,ji} = g_{ji} S(a_{ji} y_j + b_{ji})(e_{ji} - x_i)$  passes through a postsynaptic neuron's channels, when these open, as neurotransmitters released by the presynaptic neuron start binding to the receptors of the channels. Here,  $g_{ji}$  is the maximum conductance of the channels,  $S(a_{ji} y_j + b_{ji})$  is the probability of the channels to be open, with parameters  $a_{ji}$  and  $b_{ji}$ , and  $e_{ji}$  is the reversal potential of the channels.

Note that in a chemical synapse,  $x_i$  does not necessarily follow  $y_j$ . The synaptic current  $i_{s,ji}$  is 0 even if all channels are open, when  $e_{ji} = x_i$ . A network is an ODEs system:

$$\dot{x}_i = g_{li}(e_{li} - x_i) + \sum_{j=1}^{m+n} g_{ji} S(a_{ji} y_j + b_{ji})(e_{ji} - x_i) \quad (5)$$

These are the EECs (or LTCs [4]) of chemical synapses. They also ignore saturation aspects. Taking those into account, and rearranging terms as before, one obtains SEECs (or SLTCs):

$$\begin{aligned} \dot{x}_i &= -S(f_i) x_i + T(g_i) e_{li} \\ f_i &= g_{li} + \sum_{j=1}^{m+n} g_{ji} S(a_{ji} y_j + b_{ji}) \\ g_i &= g_{li} + \sum_{j=1}^{m+n} h_{ji} S(a_{ji} y_j + b_{ji}) \end{aligned} \quad (6)$$

As in gated RNNs [9],  $S(f_i)$  is the forget gate. However, the other gate is just  $e_{li}$ , the leaking potential. Moreover, each synapse has a different activation function (different  $a_{ji}, b_{ji}$ ).

3) *Neural Circuit Policies*: for short NCPs, consist of 4 layers [8]. The incoming information is received by the sensory neurons. This information is forwarded to interneurons, playing the role of a self-attention layer. Command neurons, which are connected in a recurrent fashion, store the state of the NCP, and use this together with the attention information, to compute actions. These are forwarded to the motor neurons. The NCP design aims to create a sparse architecture with interpretable dynamics. In contrast, commonly used fully-connected RNN architectures, do not distinguish layers of neurons, and neurons are connected all-to-all.

### III. RELATED WORK

Imitation learning is about training the agents to learn the connections between perceptions and actions, by observing expert demonstrations of a task [10]. Its application extends across various domains, including robotic manipulation [14], locomotion [15], and autonomous navigation, both on the ground [5], [16] and in aerial environments [17], [6], [7].

Autonomous lane-keeping solutions include end-to-end convolutional neural networks (ConvNet) [18], [19], and more advanced solutions of ConvNet heads stacked with RNNs to consider previous information during decision making [5]. Here, we use the latter variant, as similarly, we need a ConvNet head to preprocess the raw visual input which will be fed into recurrent chemical or electrical models.

The CT-RNN model of electrical synapses is a commonly used neural model, that has been applied in walking of a six-legged agent [20], robot-to-human handover [21], balancing a biped robot [22] and swarm robotics [23]. Recently, the LTC model of chemical synapses has demonstrated the ability to autonomously navigate both cars [5] and drones [6], [7].

An NCP wiring architecture built up from 19 LTC neurons, demonstrated better robustness compared to an end-to-end ConvNet, a fully-connected 64-neuron long-short term memory (LSTM) [24] and a fully-connected 64-neuron CT-RNN [5]. As this paper investigates autonomous lane-keeping too, our setup includes the two bio-inspired models from the mentioned paper: LTC neurons in a 19-neuron NCP architecture and 64 CT-RNNs all-to-all connected.

To the best of our knowledge, this paper is the first one to systematically explore the role of chemical and electrical synapses, in the same types of architectures. This combination enables us to provide a response to how much the architectural or synaptic model, respectively, has an impact on the network's accuracy, robustness, and interpretability.

TABLE I: Parameters of examined networks. The 19-neuron NCP has 12 interneurons, 6 command neurons and 1 motor neuron. The 64-neuron NCP is proportionally larger, with 42 interneurons, 21 command neurons and 1 motor neuron.

ID	Neural model	Network arch.	#Neuron	#Synapse	#Param
1	LTC	NCP	19	444	1833
2	LTC	fully	19	1577	6365
3	CT-RNN	NCP	19	444	482
4	CT-RNN	NCP	64	1700	1828
5	CT-RNN	fully	64	8192	8320

#### IV. METHODS

We thoroughly analyze the impact of the architecture and the synaptic models by conducting two types of experiments:

1) *Fix the network architecture*: In other words, change the synaptic model while keeping the number of neurons and their connections. This may result in an unfair performance comparison, because for the same network architecture, the simpler synaptic model used by CT-RNNs, offers much less trainable parameters. At the same time, the activity of the corresponding neurons is straightforward to compare.

2) *Fix the number of network parameters*: In other words, change the network architecture, such that it results in the same number of trainable parameters. This way, the networks have an equal learning capacity. However, as CT-RNNs require a much bigger network architecture, they may lose the nice properties of interpretability and compactness. There is a plain trade-off between using the same network architecture, or the same number of trainable parameters.

As network bases, we use the LTC-NCP with 19 neurons, and the fully-connected version of CT-RNN with 64 neurons, as in [5]. To close the knowledge gap between these two, we analyze the CT-RNNs in NCP architecture, with two sizes: 19 neurons (12 inter-, 6 command and 1 motor neuron), which corresponds to the type (1) experiment, and 64 neurons (42 inter-, 21 command and 1 motor neuron), which corresponds to the type (2) experiment, with approximately 1800 trainable parameters. Additionally, we include the LTCs in a fully-connected architecture, for a thorough analysis of the effect of the architecture. These setups are summarized in Table I.

##### A. Evaluation Metrics

We use several evaluation metrics to assess driving performance. They include prediction error, driving smoothness, and crash likelihood in closed-loop simulation on unknown trajectories. For interpretability, the metrics of interest are the network attention, and the changes in neuron activity during driving. In order to evaluate network robustness, in noisy conditions not present during training, we compare the attention maps and the driving ability, using the VISTA simulator. Our evaluation seeks to assess the importance of the synaptic model and the network architecture.

1) *Training and Validation Loss*: We calculate the mean squared error (MSE) between the predicted road curvature  $\hat{y}_i$  and the ground-truth  $y_i$  of the validation set. Previous

work [25] pointed out that MSE alone, may not show the full driving quality. We thus use a second metric, weighting MSE values with the road steepness, at each time step.

2) *Crash-Likelihood*: We aim to assess the performance of the networks in a closed-loop setting, where the next camera input is influenced by the previous steering decision. To this end, we measure the crash likelihood during multiple test runs. In the simulator, a crash is counted, if the car leaves the road. To test the models in even more extreme conditions, we add extra Gaussian noise to the camera input, which was not present as a data augmentation technique in the training process. First, we use mean  $\mu=0$  and variance  $\sigma^2=0.1$ , to see how they tackle this condition, and then we run experiments with doubled noise variance  $\sigma^2=0.2$ .

3) *Trajectory Smoothness*: To measure the smoothness of the driving trajectory, we compute the Lipschitz constant of the prediction over a test episode. A function  $f$  is called Lipschitz continuous with Lipschitz constant  $L \geq 0$  if it satisfies:  $|f(x_1) - f(x_2)| \leq L|x_1 - x_2|$ . If  $|x_1 - x_2|$  is the time passed between two timesteps, then  $|f(x_1) - f(x_2)|$  is the corresponding change in the prediction. The value of  $L$  measures the maximum rate of change in the network’s output. Thus smaller values indicate smoother driving.

4) *Saliency Maps and Structural-Similarity Index*: Getting insights into how a neural network makes its decision is very important in end-to-end autonomous systems. A possible approach is to analyze where the neural network pays the most attention during decision-making, that is, find out which are the pixels in the input image that have the greatest impact. We observe the distinct evolution of the identical ConvNet architecture across all the models after training. To visualize the network’s attention, we utilize the VisualBackprop algorithm [26], on our ConvNet heads. The main idea behind this method is, that the most relevant information comes from the feature maps of the deepest layers, which we bring backward layer by layer. This goes repeatedly, until the first layer, after which we take the absolute value of the attention, because the image standardization layer shifted the pixel values, and both positive and negative values should be considered with equal importance. To make one extra step, we aim to see the robustness of their learned ConvNet heads, in the presence of input noise, by comparing their saliency maps under normal and additional noise pairwise. The structural-similarity index (SSIM) [27], compares luminance, contrast, and structure of two signals, respectively. Signals are in our context images. In Figure 2 we show how to compute the SSIM metric.

5) *Neural Activity*: Another way to interpret, that is, to get insights into the network behavior, besides saliency maps, is to examine neural activity over time. Particularly focusing on whether the neural network’s individual neurons exhibit correlated activity with straight road segments, and turns, along the driving path. This is achieved by extracting the driving trajectory from the simulator, as seen from a bird’s-eye view, and then projecting the output of each neuron onto the car’s trajectory, separately. To measure the similarity be-

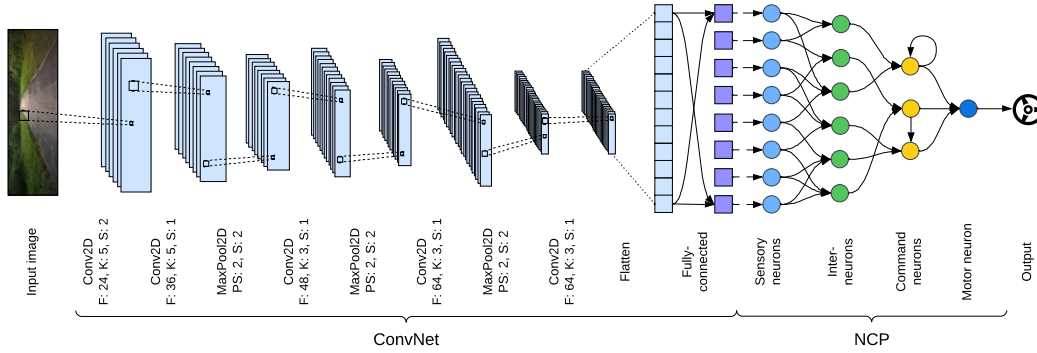


Fig. 1: The 4-layer NCP stacked with a convolutional head. The ConvNet extracts information from the input image which then serves as sensory input for the NCP. In a fully connected architecture, we replace the NCP with all-to-all wirings.

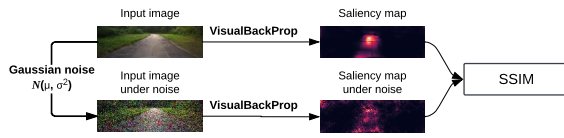


Fig. 2: Structural-similarity index (SSIM) between saliency maps, under normal and Gaussian-noise conditions.



Fig. 3: Examples from the dataset: summer and winter

tween the road curvature and the neural activity, we compute the cross-correlation of these signals. As we are interested in both negative and positive correlations, compared to no correlation, we take the absolute value of the correlations, and average them across all neurons in the network.

### B. Experimental Setup

We evaluate the performance of the NCP and all-to-all architectures, with either electrical or chemical synapses, respectively, within an end-to-end autonomous driving scenario. Here the agent receives the front camera input of the car, and predicts the road curvature, which corresponds to the steering action of the lane-keeping task. Our dataset consists of human driving recordings, under different weather conditions [28]. Figure 3 shows samples from this dataset.

Predicting the road curvature instead of directly predicting the steering angle is beneficial because the curvature value is independent from the vehicle, while the steering angle is not. The steering angle can be computed from the curvature:

$$\alpha_t = S_v \arctan(L_v y_t) \quad (7)$$

where  $\alpha_t$  is the steering angle at timestep  $t$ ,  $S_v$  and  $L_v$  are the steering-ratio and wheelbase properties of the vehicle, respectively, and  $y_t$  is the road curvature at timestep  $t$ .

For training purposes, we create an open-loop imitation learning task, from summer and winter data, which

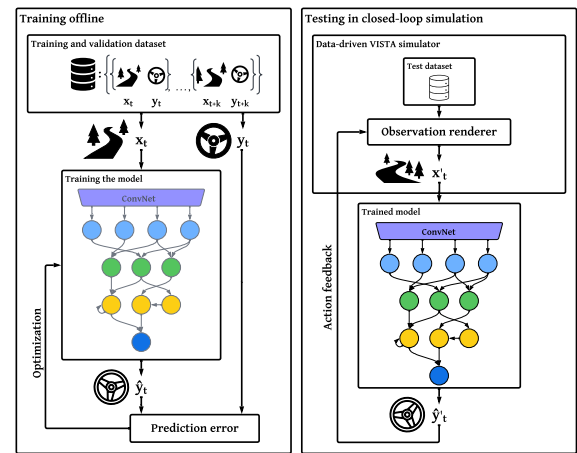


Fig. 4: Training and testing procedure.

we split into training and validation sets of 40k and 30k datapoints, respectively. The data contains input-output sequences  $\{(x_t, y_t) \mid t \in [0..T]\}$ , where  $x_t$  is the camera frame, and  $y_t$  is the road curvature at time  $t$ . The network receives the camera frame  $x_t$  at time step  $t$  and predicts the road curvature  $\hat{y}_t$ .

As data augmentation, we use the following procedure. We first adjust the brightness, by adding to the pixels a random value chosen from a uniform distribution in  $[-0.4, 0.4]$ . After this, the contrast and saturation of the image are shifted, by multiplying each RGB channel by a contrast factor, and the saturation channel in its HSV representation by a saturation factor, respectively. These random values are taken from a uniform distribution in  $[0.6, 1.4]$ . In addition to this, the training data is augmented by Guided Policy Learning [29] using the VISTA simulator [30], [11]. This places the car off-orientation on the track, and the model has to learn to recover the car from these quite challenging positions.

The model is trained using the mean squared error (MSE) of the prediction  $\hat{y}_t$  and the label  $y_t$ . This is offline learning: the prediction of the network does not influence the next input image of the sequence. This can be seen on the left-hand side in Figure 4. We run the training on a sequence of

TABLE II: Prediction error in the validation set, for 5 trained models, per model-type. Smaller error is better performance.

Model	MSE	Weighted MSE
LTC NCP-19	0.173 $\pm$ 0.021	<b>0.011 <math>\pm</math> 0.001</b>
LTC fully-19	0.200 $\pm$ 0.027	<b>0.010 <math>\pm</math> 0.001</b>
CT-RNN NCP-19	0.174 $\pm$ 0.012	0.015 $\pm$ 0.002
CT-RNN NCP-64	<b>0.163 <math>\pm</math> 0.017</b>	<b>0.012 <math>\pm</math> 0.002</b>
CT-RNN fully-64	0.219 $\pm$ 0.015	0.018 $\pm$ 0.002

input images of length 32, with a batch size of 32, with a cosine learning-rate decay schedule with an initial value of  $10^{-3}$  over 100 epochs. We apply the AdamW optimizer [31] which in our case decouples the weight decay with the value of  $10^{-6}$  from the learning rate. The best epoch is determined by the highest validation loss, without early stopping, in order to achieve a better generalization [32].

After training the models, we test them in the data-driven VISTA simulation platform, which synthesizes new possible trajectories from the dataset.  $x'_t$  represents the generated input frame and  $\hat{y}'_t$  the prediction of the model in Figure 4. This way we can evaluate the generalization performance of the models, in a closed-loop environment, where the prediction of the current time step affects the camera input of the next time step. This approach is illustrated on the right-hand side in Figure 4.

## V. RESULTS

### A. MSE in Open-Loop

Table II compares the neural networks in terms of MSE and Weighted MSE in the open-loop setting. We observe that the CT-RNN NCP-64 model exhibited the lowest MSE error, indicating the best fit to the data compared to the other models. By assigning higher weights to predictions in steeper curves, LTC-based models showed better performance to it by achieving the lowest Weighted MSE values.

### B. Crash-Likelihood in Closed-Loop

We tested the networks for 125 episodes for each season and calculated how likely they drive off the road over these episodes. We found that the chemical-synapse-based LTCs independently of the architecture, and the electrical-synapse-based CT-RNNs in NCP wiring architectures drive without crash in the closed-loop setting using the VISTA simulator. Moreover, they are robust enough, to tackle the challenge of receiving input images with 0-mean and 0.1-variance Gaussian noise. In contrast, the fully-connected CT-RNN network crashes without any input disturbances during winter and under additional noise in both summer and winter tests.

We also aimed to push these models to their limits, and observed that by doubling the noise variance, the models are not able to avoid the crashes anymore. The trend of the crash likelihood shows that the 4-layer structure of the NCP brings an advantage compared to the fully-connected variant.

Note that some of the evaluation metrics of the CT-RNN in fully-connected architecture could be not relevant, as this was not able to drive without crashes under normal circumstances during winter, and under 0.1-variance Gaussian noise in both seasons. However, we included them in our full analysis.

TABLE III: Crash likelihood in closed-loop simulation, under noise-free conditions, and under Gaussian noise with mean of 0 and variance of 0.1 and 0.2. Smaller is better. S is summer, W is winter. 125 runs per model in each condition.

Model	Noise-free		$\sigma^2 = 0.1$		$\sigma^2 = 0.2$	
	S	W	S	W	S	W
LTC NCP-19	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	0.11	<b>0.01</b>
LTC fully-19	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	0.18	0.02
CT-RNN NCP-19	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	0.20	0.08
CT-RNN NCP-64	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.05</b>	0.03
CT-RNN fully-64	<b>0.0</b>	0.02	0.09	0.06	0.68	0.25




Fig. 5: Summer-winter input images in (a), and their saliency maps of LTC NCP-19 in (b), LTC fully-19 in (c), CT-RNN NCP-19 in (d), CT-RNN NCP-64 in (e) and CT-RNN fully-64 in (f). Observe that an NCP architecture makes the models more focused compared to their fully connected version. The LTC NCP-19 keeps its attention on the road and its contours in summer, and on the road contours in winter. It ignores the irrelevant information from the border of the image.

### C. Trajectory Smoothness

Table IV provides the Lipschitz constants for networks in summer and winter contexts, both in noise-free conditions and under the influence of 0-mean and 0.1-variance noise. Observe that LTC networks generally exhibit smoother driving performance compared to CT-RNNs, in most scenarios. Moreover, when employing the NCP architecture, LTCs achieve an exceptionally smooth trajectory.

### D. Network Interpretability

1) *Saliency maps*: Figure 5 shows the saliency maps of the networks considered, for both summer and winter. We found that the networks extract information differently in the two seasons: the attention goes to the road itself in summer, and the contour (the side) of the road in winter. Intuitively, the areas of the input image containing the most relevant steering information should be highlighted: the road, and its contours, without the unnecessary environmental features which occur outside of the road. Note that the NCPs focus on the relevant parts, alone, while the fully connected networks also consider irrelevant areas of the input frames.

2) *Neural Activity*: When considering the neural activity during driving, we observed that in NCP architectures, neural activities change their dynamics based on the road curvature, regardless of the synaptic type. This behavior holds for LTCs in the fully-connected architecture, too, which means that this interpretability feature holds regardless of the network architecture, for chemical synapses. However, the electrical-synapse-based CT-RNN does not have this nice property without the NCP structure. Its all-to-all wiring shows no

TABLE IV: Lipschitz constant of the prediction in the first 4 columns. Smaller is better. Similarity of the trajectory between normal conditions and under Gaussian noise variance of 0.1 in the last 2 columns. Higher is better. LTC-based networks can maintain a smooth driving trajectory in the presence of noise better than CT-RNN networks. Using an NCP architecture makes the neural network more resilient. Note that even though the CT-RNN fully-64 results show good characteristics, this model drives off the road under this amount of noise, while the others can drive without crashes.

Model	Lipschitz constant of the output				Similarity of the trajectory	
	Noise-free		Under noise $N(\mu = 0, \sigma^2 = 0.1)$		Under noise $N(\mu = 0, \sigma^2 = 0.1)$	
	Summer	Winter	Summer	Winter	Summer	Winter
LTC NCP-19	0.0081 $\pm$ 0.0028	<b>0.0087 <math>\pm</math> 0.0031</b>	<b>0.0115 <math>\pm</math> 0.0028</b>	<b>0.0113 <math>\pm</math> 0.0036</b>	<b>0.971 <math>\pm</math> 0.043</b>	<b>0.980 <math>\pm</math> 0.027</b>
LTC fully-19	<b>0.0076 <math>\pm</math> 0.0022</b>	0.0101 $\pm$ 0.0045	0.0124 $\pm$ 0.0028	0.0126 $\pm$ 0.0041	0.967 $\pm$ 0.046	0.976 $\pm$ 0.030
CT-RNN NCP-19	0.0107 $\pm$ 0.0097	0.0150 $\pm$ 0.0098	0.0143 $\pm$ 0.0093	0.0172 $\pm$ 0.0092	0.964 $\pm$ 0.054	0.968 $\pm$ 0.047
CT-RNN NCP-64	0.0109 $\pm$ 0.0067	0.0124 $\pm$ 0.0070	0.0124 $\pm$ 0.0060	0.0148 $\pm$ 0.0070	0.962 $\pm$ 0.057	0.967 $\pm$ 0.047
CT-RNN fully-64	<b>0.0076 <math>\pm</math> 0.0029</b>	0.0107 $\pm$ 0.0089	0.0126 $\pm$ 0.0033	0.0122 $\pm$ 0.0049	0.911 $\pm$ 0.171	0.970 $\pm$ 0.040




Fig. 6: Neural activity. (a) Road trajectory from above, arrows indicate the driving direction (b)-(f) Models in the same order as before. The NCP architecture brings interpretability to CT-RNNs as the changes in neural activity correspond to the road’s trajectory, as illustrated in (d)-(e). LTC in all-to-all architecture is still interpretable in this sense, as seen in (c).

TABLE V: Correlation between the neural activity and the driving trajectory. Higher is better in this case. The NCP architecture brings the CT-RNN model much closer to the LTC-type networks, in terms of interpretable neural activity.

Model	Summer	Winter
LTC NCP-19	<b>0.662 <math>\pm</math> 0.290</b>	0.636 $\pm$ 0.279
LTC fully-19	<b>0.661 <math>\pm</math> 0.307</b>	<b>0.698 <math>\pm</math> 0.270</b>
CT-RNN NCP-19	0.649 $\pm$ 0.316	0.630 $\pm$ 0.293
CT-RNN NCP-64	0.568 $\pm$ 0.313	0.524 $\pm$ 0.282
CT-RNN fully-64	0.395 $\pm$ 0.282	0.321 $\pm$ 0.254

patterns of different activation during turns and straight sections of the road. For illustration, Figure 6 shows the activity of one of the neurons in the examined networks. In addition, we also computed the average correlation between the neural activity of each neuron, and the driving trajectory. The results in Table V clearly show that the LTC neurons have the best interpretability in terms of network dynamics, irrespective of the architecture used. In contrast, for the classic CT-RNN models, the NCP architecture gets around 1.5-2 times increased activity correlation compared to the all-to-all one. Hence, the NCP architecture brings CT-RNNs very close to the chemical-synapse-based LTC model.

### E. Network Robustness

Finally, we measured how much the saliency maps of the various networks, changed in the presence of a noisy input. An SSIM value close to 1 indicates a more robust attention mechanism, meaning that the additional noise has less of an influence on the convolutional head of these networks. In Figure 7 we present our results. These show that the LTC model in the NCP wiring architecture is the most robust: it extracts nearly identical information from the input images




Fig. 7: The structural-similarity index (SSIM) for the neural networks considered, under Gaussian noise with mean 0, and variance 0.1 and 0.2, in summer on the left-hand side and winter on the right-hand side. Higher is better.

even when noise is present.

In order to evaluate how well the five models considered can maintain their driving trajectory, we also calculated the similarity between their trajectories with and without noise. These results are reported and discussed in Table IV.

## VI. CONCLUSIONS

In this paper, we showed how neuroscience models of electrical and chemical synapses result in continuous-time recurrent neural networks. These bio-inspired neural networks hold the potential for interpretable solutions, which are highly sought after in robotics systems.

We investigated the impact of synaptic choices in the architectures of sparse and fully-connected wirings. Through extensive experiments conducted in a photorealistic autonomous driving simulator, we evaluated their performance, interpretability, and robustness under diverse conditions and in the presence of noise.

Our findings revealed that the choice of wiring architecture had a significant impact on the networks’ performance. Specifically, the sparse NCP wirings were more effective with both synaptic models. This effectiveness arises from the specialization of layers: interneurons can learn correlations of the sensory input, and recurrent command neurons process it further before sending signals to the motor neurons. Concerning the synaptic models themselves, our research demonstrated that learning with chemical synapses makes a difference. They not only led to a more robust behavior but also resulted in better interpretability and significantly smoother driving dynamics when compared to electrical synapses.

## REFERENCES

- [1] R. D. Beer and J. C. Gallagher, "Evolving dynamical neural networks for adaptive behavior," *Adaptive behavior*, vol. 1, no. 1, pp. 91–122, 1992.
- [2] K. Funahashi and Y. Nakamura, "Approximation of dynamical systems by continuous time recurrent neural networks," *Neural Networks*, vol. 6, pp. 801–806, 1993.
- [3] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Advances in Neural Information Processing Systems*, vol. 31, pp. 1–13, Curran Associates, Inc., 2018.
- [4] R. M. Hasani, M. Lechner, A. Amini, D. Rus, and R. Grosu, "Liquid time-constant networks," in *AAAI Conference on Artificial Intelligence*, 2020.
- [5] M. Lechner, R. M. Hasani, A. Amini, T. A. Henzinger, D. Rus, and R. Grosu, "Neural circuit policies enabling auditable autonomy," *Nature Machine Intelligence*, vol. 2, pp. 642–652, 2020.
- [6] C. J. Vorbach, R. M. Hasani, A. Amini, M. Lechner, and D. Rus, "Causal navigation by continuous-time neural networks," in *Neural Information Processing Systems*, 2021.
- [7] M. Chahine, R. Hasani, P. Kao, A. Ray, R. Shubert, M. Lechner, A. Amini, and D. Rus, "Robust flight navigation out of distribution with liquid neural networks," *Science Robotics*, vol. 8, no. 77, p. eadc8892, 2023.
- [8] M. Lechner, R. Hasani, M. Zimmer, T. A. Henzinger, and R. Grosu, "Designing worm-inspired neural networks for interpretable robotic control," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 87–94, 2019.
- [9] G.-B. Zhou, J. Wu, C.-L. Zhang, and Z.-H. Zhou, "Minimal gated unit for recurrent neural networks," 2016.
- [10] S. Schaal, "Is imitation learning the route to humanoid robots?," *Trends in Cognitive Sciences*, vol. 3, pp. 233–242, 1999.
- [11] A. Amini, T.-H. Wang, I. Gilitschenski, W. Schwarting, Z. Liu, S. Han, S. Karaman, and D. Rus, "Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2419–2426, IEEE, 2022.
- [12] S. R. Wicks, C. J. Roehrig, and C. H. Rankin, "A dynamic network simulation of the nematode tap withdrawal circuit: predictions concerning synaptic function using behavioral criteria," *Journal of Neuroscience*, vol. 16, no. 12, pp. 4017–4031, 1996.
- [13] E. R. Kandel, J. H. Schwartz, T. M. Jessell, S. Siegelbaum, A. J. Hudspeth, S. Mack, et al., *Principles of neural science*, vol. 4. McGraw-hill New York, 2000.
- [14] E. Johns, "Coarse-to-fine imitation learning: Robot manipulation from a single demonstration," in *2021 IEEE international conference on robotics and automation (ICRA)*, pp. 4613–4619, IEEE, 2021.
- [15] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," *arXiv preprint arXiv:2004.00784*, 2020.
- [16] H. Karnan, G. Warnell, X. Xiao, and P. Stone, "Voila: Visual-observation-only imitation learning for autonomous navigation," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2497–2503, IEEE, 2022.
- [17] Y. Fan, S. Chu, W. Zhang, R. Song, and Y. Li, "Learn by observation: Imitation learning for drone patrolling from videos of a human navigator," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5209–5216, IEEE, 2020.
- [18] M. Bojarski, D. W. del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, "End to end learning for self-driving cars," *ArXiv*, vol. abs/1604.07316, 2016.
- [19] H. Xu, Y. Gao, F. Yu, and T. Darrell, "End-to-end learning of driving models from large-scale video datasets," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3530–3538, 2016.
- [20] R. D. Beer, "A dynamical systems perspective on agent-environment interaction," *Artificial intelligence*, vol. 72, no. 1-2, pp. 173–215, 1995.
- [21] H. Tang, B. H. Tan, and R. Yan, "Robot-to-human handover with obstacle avoidance via continuous time recurrent neural network," in *2016 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1204–1211, 2016.
- [22] P. Hénaff, V. Scesa, F. B. Ouedzou, and O. Bruneau, "Real time implementation of ctrnn and bptt algorithm to learn on-line biped robot balance: Experiments on the standing posture," *ArXiv*, vol. abs/2011.06910, 2020.
- [23] R. Sendra-Arranz and Á. Gutiérrez, "Emergence of communication through artificial evolution in an orientation consensus task in swarm robotics," in *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pp. 515–526, Springer, 2023.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] F. Codevilla, A. M. Lopez, V. Koltun, and A. Dosovitskiy, "On offline evaluation of vision-based driving models," in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 236–251, 2018.
- [26] M. Bojarski, A. Choromańska, K. Choromanski, B. Firner, L. D. Jackel, U. Muller, and K. Zieba, "Visualbackprop: visualizing cnns for autonomous driving," *ArXiv*, vol. abs/1611.05418, 2016.
- [27] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [28] M. Lechner, R. Hasani, A. Amini, T.-H. Wang, T. A. Henzinger, and D. Rus, "Are all vision models created equal? a study of the open-loop to closed-loop causality gap," *arXiv preprint arXiv:2210.04303*, 2022.
- [29] S. Levine and V. Koltun, "Guided policy search," in *Proceedings of the 30th International Conference on Machine Learning (S. Dasgupta and D. McAllester, eds.)*, vol. 28 of *Proceedings of Machine Learning Research*, (Atlanta, Georgia, USA), pp. 1–9, PMLR, 17–19 Jun 2013.
- [30] A. Amini, I. Gilitschenski, J. Phillips, J. Moseyko, R. Banerjee, S. Karaman, and D. Rus, "Learning robust control policies for end-to-end autonomous driving from data-driven simulation," *IEEE Robotics and Automation Letters*, vol. 5, pp. 1143–1150, 2020.
- [31] I. Loshchilov and F. Hutter, "Fixing weight decay regularization in adam," *ArXiv*, vol. abs/1711.05101, 2017.
- [32] A. Power, Y. Burda, H. Edwards, I. Babuschkin, and V. Misra, "Grokking: Generalization beyond overfitting on small algorithmic datasets," *ArXiv*, vol. abs/2201.02177, 2022.