Towards endowing intelligent cars with the ability to learn the routines of multiple drivers: a dynamic neural field model *

Weronika Wojtak[1,2], Flora Ferreira[1], Pedro Guimarães[1,2], Paulo Barbosa[1,2], Sérgio Monteiro[2], Wolfram Erlhagen[1], and Estela Bicho[2]

[1] Research Centre of Mathematics, University of Minho, 4800-058 Guimarães, Portugal
{fjferreira,wolfram.erlhagen}@math.uminho.pt,
[2] Research Centre Algoritmi, University of Minho, 4800-058 Guimarães, Portugal
{w.wojtak,estela.bicho}@dei.uminho.pt

Abstract. Driving a car is often a routine activity that includes visiting the same locations at about the same time on a certain day of the week. Here, we present a learning system based on Dynamic Neural Fields that allows an intelligent vehicle to acquire sequential and temporal information about daily driving routines. Learning is fast, implicit (no need to specify destinations in advance), continuous, and can be adapted to different temporal scales. The learned information can be recalled to predict the driver's destination intention, when to arrive at a specific location and when to leave there. Importantly, the system allows to learn and recall multiple routines corresponding to different drivers and different days. This personalized information can be used for planning the next trip for every user of a shared vehicle.

Keywords: Neurocomputational model · Learning driver routines · Space and time prediction · Dynamic neural field.

1 Introduction

Advances in technology used in the vehicle's cockpit systems have contributed to the increase of spatiotemporal information on human mobility, opening new opportunities in the research topic of human-machine interaction. For a convenient human-machine interaction, an intelligent system such as a smart vehicle must be able to learn and make decisions based on the received data. Nowadays, the Global Positioning System (GPS) is often used to guide a driver to

* The work received financial support from European Structural and Investment Funds in the FEDER component, through the Operational Competitiveness and Internationalization Programme (COMPETE 2020) and national funds, through FCT (Project **"Neurofield"**, ref POCI01-0145FEDER-031393) and ADI (Project **"Easy Ride:Experience is everything"**, ref POCI-01-0247-FEDER-039334), FCT PhD fellowship PD/BD/128183/2016, and R&D Units Project Scope: UIDB/00319/2020 and UIDB/00013/2020.

a destination that he/she enters into the system. Future vehicles might be able to predict where a specific driver wishes to go as well as the time the driver should arrive or depart [15]. Human mobility is characterized by a high level of spatial-temporal regularity [3,11,24], a tendency to visit specific locations at specific times [11,14,18,23], and a significant tendency to spend most of the time in a few locations [23]. Driving a car is typically coupled with daily routines such as going to work every weekday or routines across other temporal scales such as going to the gym on specific days of the week. Profiling the mobility routine of a specific driver can be used by intelligent navigation systems to support individual traveling by for instance warning the driver about traffic conditions that might cause a late arrival at the specific destination. Since a car might be shared by several users, it is important to keep in mind that the system should be able to make predictions according to the learned routines of different drivers. Personalization for adapting to different drivers is an important aspect in designing advanced driving assistance systems. The focus of research in this area has been thus far on systems that address safety issues by monitoring the human driving behavior and style (for recent reviews see [12,13,26]). The potential of personalization should be also taken into account in profiling mobility routines to reduce the cognitive load of the driver and ensure optimal user experience.

Several different approaches, most of them statistical models based on big data [2,21], have been proposed for predicting the next location in human mobility. Traditional Markov models work well for a specific set of behaviors but destinations need to be fixed in advance. Recently deep learning techniques have been applied for an accurate prediction of user destination [5,25]. In [5] a Long Short-Term Memory recurrent neural network is used not only to predict drivers' destinations but also respective departure times. In this framework, the destinations (places visited by a driver with significant frequency) are first identified using a clustering algorithm, and the updating of new destinations requires the learning system to be restarted. The deep learning approach in [25] is based on a division of the spatial map in grids, each with a specific ID number, and the IDs of all possible destinations are predefined at the beginning of the learning process. Although the model takes into account the temporal characteristics of the GPS trajectories, this information is not used to predict departure or arrival times.

In [8] a model based on the theoretical framework of Dynamic Neural Fields (DNFs) was proposed that is able to learn ordinal and temporal aspects of a daily driver routine. The learning is implicit (driver does not need to be asked for his/her destinations), continuous, and can be adapted to different temporal scales. The theoretical framework of DNFs has been proven to provide key brain-inspired processing mechanisms to implement working memory, decision making, and prediction in cognitive systems (e.g. [19,20]), including the learning of the temporal and ordinal properties of sequential tasks [7,10], and in recent work, the temporal integration of GPS coordinates for the identification of stop locations [9]. The central idea of DNF models is that continuous-valued information, such as for example position in space, is represented by localized activity

patterns (or bumps) in a network of recurrently connected neurons tuned to the continuous dimension. Initially driven by transient input signals, a bump becomes self-sustained due to the recurrent interactions within the network. The bump attractor thus encodes in its position the memory of a specific input value. An additional threshold accommodation dynamics ensures that the bump amplitude increases monotonically with elapsed time since input onset. In response to an entire input sequence, the field dynamics develops a multi-bump pattern with an activation gradient over neural subpopulations which carries the information about the temporal order and relative timing of the sequential events.

In this paper, we present an extension of our previous DNF model of learning daily routines [8]. The most significant advances are the integration of a long-term memory that supports the continuous up-dating and consolidation over weeks of the memory of driving experiences at a specific weekday, and the capacity to learn and distinguish the routines of different drivers. The proposed mechanism supports 1) the encoding and recall of different routines according to the driver and the day, 2) the continuous adaptation of these routines (e.g., a destination can be added or deleted without the need to restart the system), 3) the prediction of the driver's destination intention, 4) the prediction when she/he intends to arrive, and 5) the prediction how long the stay duration will be.

In what follows, we begin with a description of the DNF model and its mathematical details in Section 2. The results of model simulations of the routine learning and prediction processes based on recorded real-world GPS trajectories are discussed in Section 3. Conclusions and future work are presented in Section 4.

2 The model

2.1 Model description

Dynamic neural fields (DNFs) have been first proposed as a theoretical framework to analyze the dynamics of pattern formation in neuronal populations [1]. They have been later used to model cognitive functions such as working memory, decision making or motor planning and to synthesize these functions in autonomous robots [4]. Neural field models, formalized by integro-differential equations, represent a specific class of recurrent neural networks with a connectivity function depending on the distance in parametric space between neurons tuned to continuous metric dimensions. In the present application, the dynamic fields are spanned over the two spatial GPS coordinates latitude and longitude. The recurrent excitatory and inhibitory interactions within the network support the existence of self-stabilized bumps of activity which are initially triggered by sufficiently strong inputs. The bump position represents the memory of the GPS coordinates of a stop location of a given driver. A *stop* here stands for a visited location where the car arrived at time t_{off} and left at time t_{on}, $S = (lat, long, t_{off}, t_{on})$, where $(lat, long)$ represent the latitude and longitude coordinates corresponding to the centroid of the respective bump. Let S_i and

S_k be two stops where the index indicates different days, then $i = k$ (that is, S_i and S_k represent a stop at about the same location and time of the day) if the haversine distance between the two sets of coordinates are below a distance threshold δ_d, that is, $dist_{ik} = dist_{haversine}(lat_i, long_i, long_k, lat_k) < \delta_d$, and $|t_{off_i} - t_{off_k}| < \delta_t$, where δ_t represents a time threshold. The haversine distance [22] between two points, p_1 and p_2, with coordinates $(lat_1, long_1)$ and $(lat_2, long_2)$, respectively, is defined as

$$dist_{haversine}(p_1, p_2) = 2R\left(\sqrt{\frac{a(p_1, p_2)}{a(p_1, p_2) - 1}}\right), \tag{1}$$

where R is the Earth's radius and

$$a(p_1, p_2) = \sin^2\left(\frac{lat_2 - lat_1}{2}\right) + \cos(lat_1)\cos(lat_2)\sin^2\left(\frac{long_2 - long_1}{2}\right). \tag{2}$$

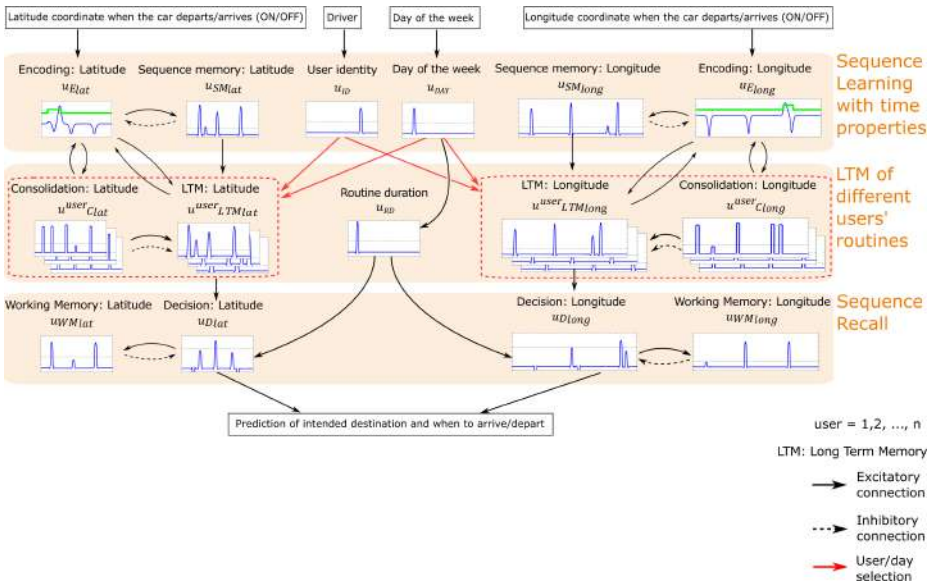

Fig. 1: Schematic view of the model architecture with several interconnected neural fields implementing sequence learning, long term memory and sequence recall. For details see the text.

Figure 1 illustrates an overview of the model architecture consisting of several interconnected neural fields. The architecture can be divided into three functionally different parts running in parallel: sequence learning, long term memory and sequence recall.

The *Sequence learning* part processes incoming GPS inputs and memorizes the stop locations of a certain driver at a specific weekday. The inputs are given by the GPS coordinates (latitude and longitude) when the vehicle is switched off or on, representing the coordinates of a destination at the time of the car's arrival or departure, respectively. We further assume that additional inputs identify the specific user and the day of the week. For consistency, we also use bumps to store this discrete information in an user identity field, u_{ID} and a weekday field, u_{DAY}. The fields are divided in several distinct regions which receive localized input in form of a Gaussian carrying the user and weekday information.

To simplify the following model description, we refer to the "arrival" signal only. At the moment when the driver switches off the car at a specific location, the GPS input drives the evolution of a bump in the latitude field $u_{E_{lat}}$ and the longitude field $u_{E_{long}}$. Each of these bumps triggers through excitatory connections the evolution of a localized activity pattern at the corresponding sites in the latitude and longitude sequence memory fields, $u_{SM_{lat}}$ and $u_{SM_{long}}$, respectively. Inhibitory feedback from $u_{SM_{lat}}$ to $u_{E_{lat}}$ (from $u_{SM_{long}}$ to $u_{E_{long}}$) destabilizes the existing bump in the encoding field. This ensures that newly arrived localized input to $u_{E_{lat}}$ ($u_{E_{long}}$) will automatically create a bump at a different field location even if the position information is repeated during the course of the day. A series of GPS coordinates at the moments when the car is switched off creates a multi-bump pattern in $u_{SM_{lat}}$ and in $u_{SM_{long}}$ with the strength of activation decreasing from bump to bump as a function of elapsed time since the start of the routine. This activation gradient over subpopulations is achieved by applying a threshold accommodation dynamics to the self-sustained activity patterns which lead to a continuous increase of the bump amplitude as a function of elapsed time [10].

The *Long Term Memory (LTM)* part stores in separate fields multiple memories of daily routines corresponding to different drivers and different days of the week. According to the information represented by a bump in the user detection and weekday fields, the corresponding consolidation fields, $u^{user}_{C_{lat}}$ and $u^{user}_{C_{long}}$, and the associated LTM fields, $u^{user}_{LTM_{lat}}$ and $u^{user}_{LTM_{long}}$, receive excitatory input from u_E and u_{SM}. The consolidation field controls whether a new stop location becomes stored in LTM as part of a daily routine or whether a memorized stop location will be forgotten since the driver has changed the routine. When a stop event occurred with a certain frequency, the accumulated localized activity in $u^{user}_{C_{lat}}$ ($u^{user}_{C_{long}}$) is above threshold. The combined excitatory input from $u^{user}_{C_{lat}}$ ($u^{user}_{C_{long}}$), $u_{E_{lat}}$ ($u_{E_{long}}$) and $u_{SM_{lat}}$ ($u_{SM_{long}}$) is then able to create a memory bump in the LTM fields. If on the other hand a bump exists in $u^{user}_{LTM_{lat}}$ ($u^{user}_{LTM_{long}}$) but the activity in $u^{user}_{C_{lat}}$ ($u^{user}_{C_{long}}$) is subthreshold due to a continuous activity decay without excitation from the encoding field $u_{E_{lat}}$ ($u_{E_{long}}$), the LTM bump becomes suppressed mediated by inhibition. There exist also excitatory connections from the memory fields back to $u_{E_{lat}}$ and $u_{E_{long}}$. This feedback excitation causes a pre-activation of neural populations in the encoding fields representing previously visited locations. Functionally, this pre-shaping mechanisms increases the robustness of the encoding process in the face of noisy and

potentially incomplete GPS inputs. The routine duration field, u_{RD}, receives excitatory input from the weekday field, u_{DAY}, at the beginning of the day. The amplitude of the resulting self-sustained activity pattern increases continuously due to applied threshold accommodation dynamics. At the end of the day, the bump amplitude represents a time span of 24 hours. The amplitude value is used to define the baseline activity level in the decision fields.

In the *Sequence Recall* phase, the stored information is used to make predictions about the driver's intended destinations and the time of arrival. The decision field $u_{D_{lat}}$ ($u_{D_{long}}$) receives the activation gradient stored in $u^{user}_{LTM_{lat}}$ ($u^{user}_{LTM_{long}}$) as subthreshold input. A continuous increase of the baseline activity in $u_{D_{lat}}$ ($u_{D_{long}}$) brings all subpopulations closer to the threshold for the evolution of a self-stabilized bump. The moment when the activity of the population with the highest pre-activation in the decision field reaches threshold is used to predict the location and arrival time of the first destination. Excitatory-inhibitory connections between associated populations in $u_{D_{lat}}$ ($u_{D_{long}}$) and the working memory field $u_{WM_{lat}}$ ($u_{WM_{long}}$) guarantee that a bump representing the coordinate of a predicted stop event evolves in the working memory field which then suppresses the suprathreshold activity in the decision field. This dynamic process continuous until the population with the lowest pre-activation has reached threshold and the prediction about the last stop event has been stored in working memory.

2.2 Model equations

The population dynamics in each field is governed by an integro-differential equation first proposed and analyzed by Amari [1]

$$\tau \frac{\partial u(x,t)}{\partial t} = -u(x,t) - h + I(x,t) + \int_{\Omega} w(|x-x'|) f(u(x',t)-\theta) \mathrm{d}x' + \epsilon^{1/2} \mathrm{d}W(x,t). \tag{3}$$

The variable $u(x,t)$ represents the activity at time t of a neuron at field position x in a spatial domain $\Omega \subset \mathbb{R}$. The constant $\tau > 0$ defines the time scale of the field dynamics. Term $I(x,t)$ represents a time-dependent, localized input centered at site x, and $h > 0$ defines the stable resting state of a field without external input. The distance-dependent connectivity function $w(|x-x'|)$ determines the interaction strength between neurons at positions x and x'. An example is a kernel of lateral inhibition type given by a Gaussian function minus a constant

$$w_{lat}(x) = A_{lat} e^{\left(-x^2/2\sigma^2_{lat}\right)} - g_{lat}, \tag{4}$$

with $A_{lat} > g_{lat} > 0$ and $\sigma_{lat} > 0$. We use the lateral inhibition kernel in the fields in which only one bump at a time should evolve (e.g., u_E and u_R) [1]. To enable stable multi-bump solutions in the memory fields, an oscillatory connectivity function [6,16] is used:

$$w_{osc}(x) = A_{osc} e^{-b|x|} (b \sin|\alpha x| + \cos(\alpha x)), \tag{5}$$

where A_{osc} controls the amplitude and parameters $b < \alpha \leq 1$ control the rate at which the oscillations decay with distance and the zero crossings of w, respectively.

The firing rate function $f(u)$ is taken here as the Heaviside step function with threshold $\theta = 0$. Finally, the additive noise term $\mathrm{d}W(x,t)$ describes the increment of a spatially dependent Wiener process with amplitude $\epsilon \ll 1$.

The activation gradients in $u_{SM_{lat}}$ and $u_{SM_{long}}$ encode the serial order of visited locations, that is, the earlier a certain location was visited, the higher is its memory bump. Since the bump amplitude increases as a function of elapsed time, the difference in bump amplitude of two successive destinations represents the temporal interval separating the two visits. To establish the activation gradient, we consider the following state-dependent dynamics [7,10]

$$\tau_h \frac{\partial h_{SM}(x,t)}{\partial t} = \kappa f(u_{SM}(x,t)) + (1 - f(u_{SM}(x,t)))(-h_{SM}(x,t) + h_{SM_0}) \quad (6)$$

where h_{SM_0} defines the baseline activation to which h_{SM} relaxes without suprathreshold activity at position x, $\kappa > 0$ is the growth rate when it is present and τ_h is time scale of the dynamics. The time window for the buildup is proportional to the total routine time, which in the present example is assumed as 24 hours. The h-level accommodation dynamics begins at $t = t_{start}$ (here, at 0:00) and ends at $t = t_{end}$ (here, at 24:00).

To ensure that only the places visited with some minimal frequency are memorized and that the places no longer visited are forgotten, the consolidation field u_C governed by equation (7) integrates the input from u_E with a simple linear dynamics. It approaches with the growth rate λ_{build} the value $a > 0$, representing the maximum level of activation in u_C. When there is no suprathreshold activity in u_E, the activity of excited neurons decays with a slower rate λ_{decay} towards the resting level $h_c < 0$:

$$\begin{aligned} \frac{\partial u_C(x,t)}{\partial t} &= \lambda_{build}(-u_C(x,t) + a) f(u_E(x,t)) \\ &\quad + \lambda_{decay}(-u_C(x,t) + h_c)(1 - f(u_E(x,t))). \end{aligned} \quad (7)$$

In the examples of user and weekday specific routines presented here, the growth and decay rates are chosen in such a way that 1) accumulated localized activity in u_C is above threshold when a certain stop point has been visited in two consecutive weeks, and that 2) the activity falls below threshold when the predicted visit has not been realized in two consecutive weeks at the specific weekday. Note that longer integration and forgetting periods could be realized by adapting the parameters λ_{build} and λ_{decay} accordingly.

A dynamic build-up of a long term memory of routines is performed in $u^{user}_{LTM_{lat}}$ ($u^{user}_{LTM_{long}}$) through excitatory connections from $u_{SM_{lat}}$ ($u_{SM_{long}}$), $u_{E_{lat}}$ ($u_{E_{long}}$) and $u_{C_{lat}}$ ($u_{C_{long}}$). The LTM fields receive the following input:

$$I_{user}(x,t) = D_{user} f(u_C(x,t)) f(u_E(x,t)) u_{SM}(x,t) \quad (8)$$

where

$$D_{user} = \max_{A_{ID}}(f(u_{ID}(x,t) - \theta)) \times \max_{A_{DAY}}(f(u_{DAY}(x,t) - \theta)), \quad user = 1, 2, \ldots, n, \tag{9}$$

and A_{ID} and A_{DAY} represent the regions in the fields u_{ID} and u_{DAY} which receive the information about the user identity and the working day, respectively. Since $D_{user} = 1$ when bumps exist in both fields and 0 otherwise, the parameter D_{user} acts as a multiplicative signal that gates the input from the encoding and sequence memory fields to the associated LTM fields.

The gradient in the LTM fields is established by the following state-dependent dynamics:

$$\begin{aligned} \tau_h \frac{\partial h_{LTM}(x,t)}{\partial t} &= (1 - f(u_{LTM}(x,t)))\,(-h_{LTM}(x,t) + h_{LTM_0}) \\ +f(u_{LTM}(x,t))f(u_C(x,t))f(u_{SM}(x,t))(-h_{LTM}(x,t) &+ h_{SM}(x,t)) \\ +f(u_{LTM}(x,t))(1 - f(u_C(x,t)))(-h_{LTM}(x,t) &+ h_{LTM_{inih}}) \end{aligned} \tag{10}$$

where $h_{LTM_0} < 0$ is the baseline activity to which h_{LTM} converges without suprathreshold activity at position x. The negative constant $h_{LTM_{inih}} < 0$ represents the stable state to which h_{LTM} converges at sites with a bump in $u_{LTM}(x,t)$ and no suprathreshold activity in $u_C(x,t)$. This ensures that the memory of a stop location that is not part of the daily routine anymore will be erased. At positions x with suprathreshold activation in the three fields $u_C(x,t)$, $u_{LTM}(x,t)$ and $u_{SM}(x,t)$, the h-value $h_{LTM}(x,t)$ tends to $h_{SM}(x,t)$. The bump amplitude in $u_{LTM}(x,t)$ thus equals the corresponding bump amplitude in $u_{SM}(x,t)$. This means that the relative timing information recorded during the last routine execution is preserved in long term memory.
The recall of the learned routine is performed by the four fields shown on bottom of Figure 1. The decision field $u_{D_{lat}}$ ($u_{D_{long}}$) receives the multi-bump patterns from $u^{user}_{LTM_{lat}}$ ($u^{user}_{LTM_{long}}$) as subthreshold input

$$I_{user}(x) = \sum_{user=1}^{n} D_{user} u^{user}_{LTM}(x). \tag{11}$$

The coefficient D_{user} given by (9) ensures again that only the gradient that corresponds to the current driver and day of the week contributes to this input.

To recall the order and timing information stored in the activation gradients, we apply a linear dynamics for the resting level h_D in $u_{D_{lat}}$ ($u_{D_{long}}$) which brings all pre-activated subpopulations closer to the threshold 0:

$$\frac{\mathrm{d}h_D(t)}{\mathrm{d}t} = \kappa_D, \quad h_D(t_0) = h_{D_0} < 0, \tag{12}$$

where κ_D defines the constant growth rate.

The initial value h_{D_0} is chosen equal to the amplitude of the bump in the duration field, u_{RD}, from the last routine execution. Since the amplitude of this

bump represents a routine time span of 24 hours, the representation of the stop events reaches the threshold 0 exactly at the predicted event time. The initial resting state could be chosen slightly larger to ensure that the diver assistant anticipates (e.g., by $t_a = 10$ minutes) the true arrival time. Note that we have proposed in our previous work a learning mechanism that allows the system to autonomously adapt its internal event timing encoded in the bump amplitudes based on environmental feedback [10].

Numerical simulations of the model were performed in MATLAB using a forward Euler method. To compute the spatial convolution of w and f we employ a fast Fourier transform (FFT), using MATLAB's in-built functions `fft` and `ifft` to perform the Fourier transform and the inverse Fourier transform, respectively.

3 Results

To test the proposed model, we consider a previously recorded dataset consisting of 11 consecutive weeks of a real-world driving scenario. It represents the routines of two different drivers from the Portuguese city Guimarães sharing the same car. To identify distinct stop locations in the dataset, we apply threshold values for the location radius ($\delta_d = 200$ meters) and the stay duration ($\delta_t = 30$ minutes). Furthermore, we assume that a stop event is part of daily routine if it happens at least twice in consecutive weeks and that the location is forgotten when a prediction is not confirmed twice in a row. To generate predictions, a routine for each day and each driver has to be executed at least three times. To directly compare the real with the predicted arrival times, we do not apply any forecast period, that is, a perfect prediction matches the exact arrival time.

As an example of a daily routine of one of the drivers, Figure 2(A) shows a map with a series of five destinations where the car was switched off, one place (the school) was visited twice. The memories of the GPS coordinates, latitude and longitude, are represented in the activation gradients in $u_{SM_{lat}}$ (left) and $u_{SM_{long}}$ (top), respectively. Supposing that this daily routine does not change in space and time, running the routine three times results in long term memories (not shown) reflecting the order and the timing of the arrival/departure destinations. Fig. 2(B) compares the exact point in time in which the driver arrives at and departures from a specific destination (vertical lines) with the time course of the maximal activation of the corresponding population representation in the decision fields, where $u_{D_{long}}$ is taken as example. As can be seen when considering the reaching of the threshold as read-out time, there is a perfect match. The temporal difference between successive time courses indicates the interval timing between successive arrivals or departures. The information about the stay duration at a specific destination can be extracted by comparing the respective time courses of arrival and departure. In the specific example considered in Fig. 2(B), the stay duration predictions at stops S_1 to S_5 are: S_1: 21 minutes, S_2: 7 hours and 54 minutes, S_3: 20 minutes, S_4: 60 minutes and S_5: 24 minutes.

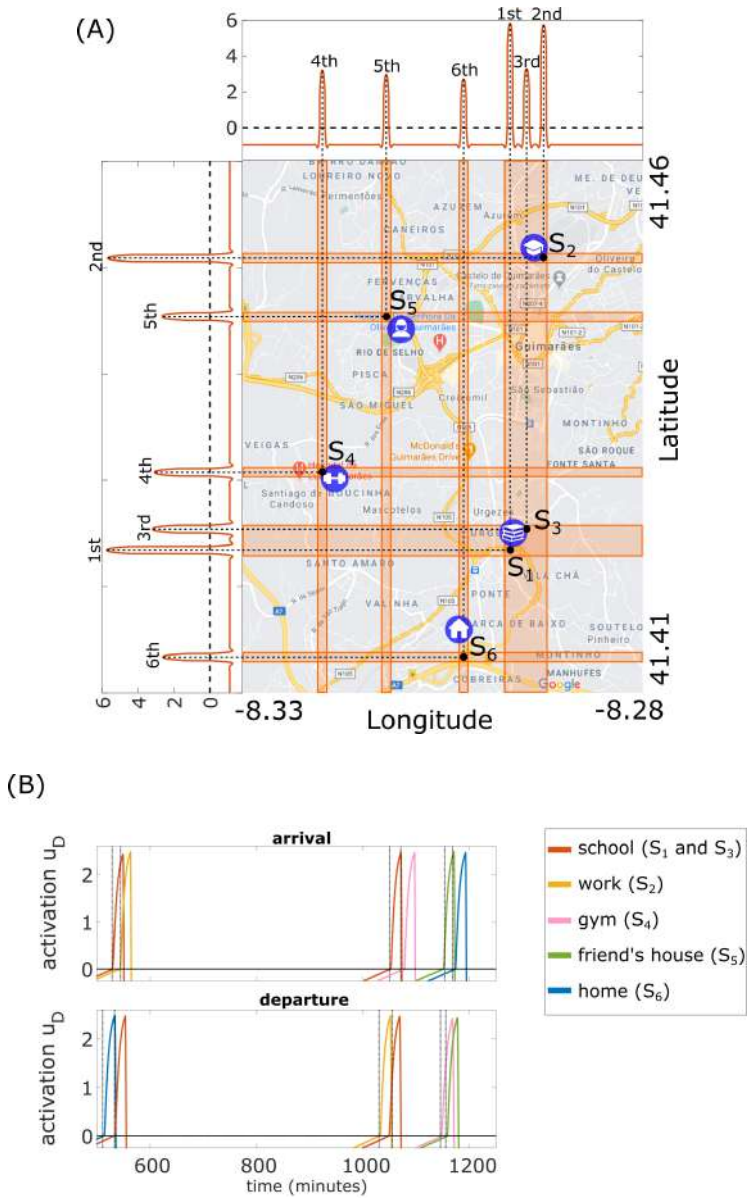

Fig. 2: (A) Map of part of Portugal generated from Google Maps showing a sequence of arrivals of Driver A and the respective activation gradients in the memory fields $u_{SM_{lat}}$ (left) and $u_{SM_{long}}$ (top). (B) Time courses of the activation at the bump centers in the decision field $u_{D_{long}}$.

As an example to compare daily routines of the two drivers, we consider 11 Monday trajectories of driver A and 8 Friday trajectories of driver B. Looking at the time t_{off} when the driver has switched off the car, Table 1 compares the real and predicted arrival times of each stop for both driver routines. Driver A stops at locations S_1, S_3, S_4 and S_5 on all Mondays. From the third Monday on, the stops are correctly predicted and the foreseen arrival times over the weeks closely follow the observed variations in timing of the driver trajectories. For stop locations such as S_1 with a relatively small variation of real stop times, the differences between predicted and realized arrival time is small (Figure 3(A)). Adding a fixed forecast period that takes into account this variability would guarantee that a driver assistant anticipates in all cases the real stop event. Stop location S_6 is an example in which the arrival time appears to be significantly delayed in one week compared to the habitual timing pattern (compare Mondays 4 and 5 in Figure 3(C)). Since the arrival delay does not represent a consistent change in the routine and the timing of the current prediction reflects the observed timing of the preceding routine execution, a relatively large prediction delay manifests in the following week. Stop S_2 did not occur on all Mondays, but this stop location was not completely forgotten since its memory is refreshed in the week directly following a trajectory without visit (e.g., weeks 6 and 7 in Figure 3(B)). Driver A stopped at S_5 only twice in different weeks. This stop event thus does not enter the long term routine memory and no prediction is made. Driver B stops at locations S_1, S_2, S_3, and S_5 on all Fridays. From the third Friday on, these four stops are correctly predicted with a timing pattern that closely matches the mean and standard deviation of the observed pattern. Implementing a short forecast period in the field dynamics would again ensure that the prediction anticipates in all cases the real arrival time. Stop S_4 happened only on the first three Fridays (Figure 3(D)). From the third to the fifth Friday the systems incorrectly predicts this location as being part of a routine, and in the sixth week it appears to be forgotten.

Table 1: Real and predicted times in minutes of two different day routines from two different drivers.

Monday routine of driver A (11 weeks)					Friday routine of driver B (8 weeks)				
Stops	NS	Real time	NP	Predicted time	Stops	NS	Real time	NP	Predicted time
S_1	11	526.5 (±1.81)	9	525.8 (±1.97)	S_1	8	520.0 (±4.11)	6	518.3 (±5.16)
S_2	7	542.1 (±2.19)	9	541.1 (±2.66)	S_2	8	539.5 (±5.88)	6	537.0 (±6.95)
S_3	11	550.1 (±6.45)	9	549.2 (±7.18)	S_3	8	1116.3 (±8.40)	6	1117.0 (±9.67)
S_4	11	1113.1 (±2.07)	9	1112.8 (±2.87)	S_4	3	1132.0 (±3.61)	3	1133.4 (±3.64)
S_5	2	1129.5 (±12.02)	0	–	S_5	8	1261.6 (±5.21)	6	1262.5 (±5.48)
S_6	11	1134.0 (±7.35)	9	1134.0 (±7.67)	–				

Data is presented as mean (± standard deviation). NS and NP are the number of weeks that the location at about the same period of the day was visited and predicted, respectively.

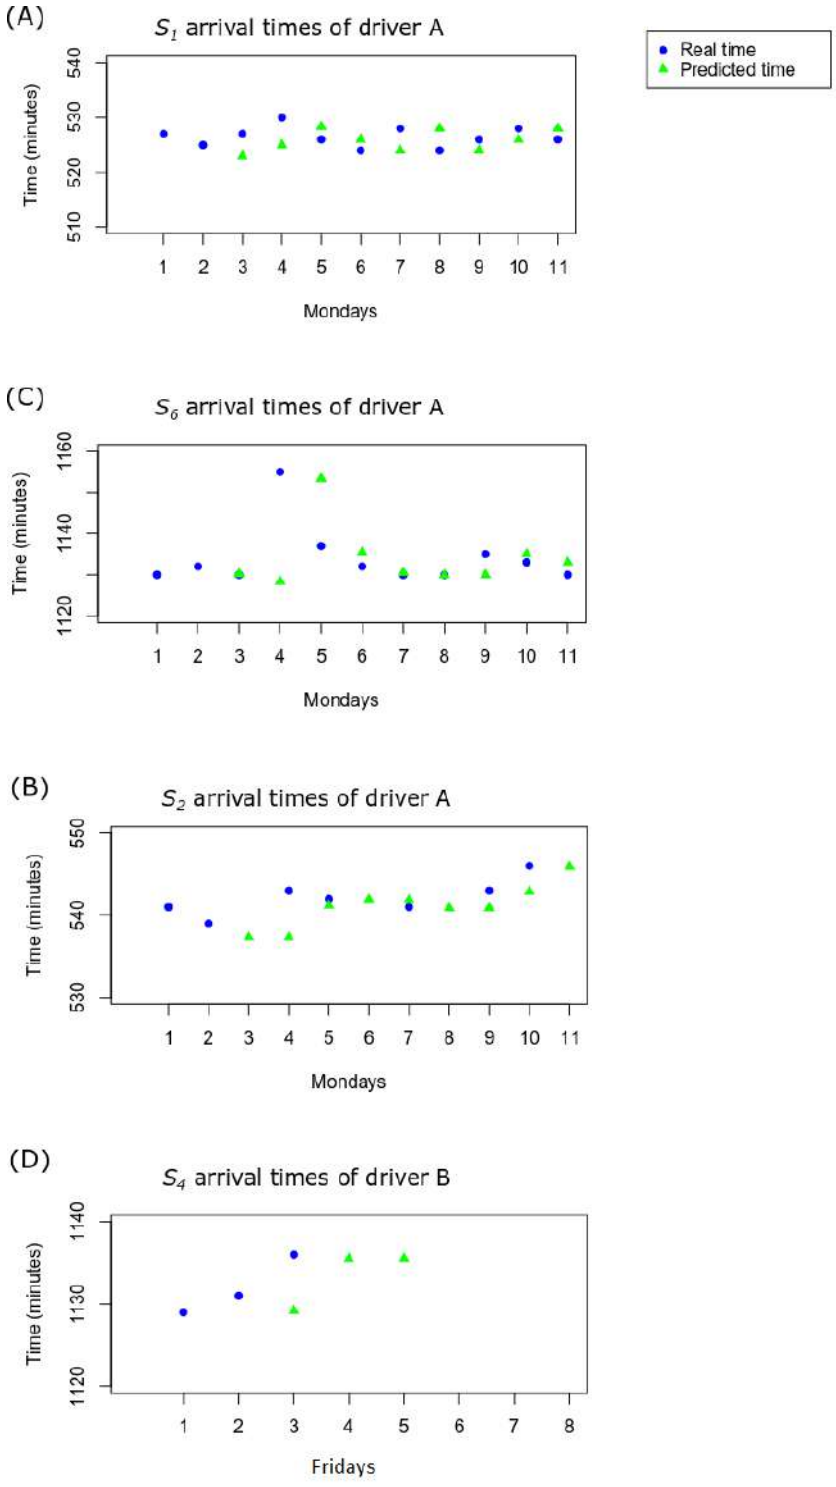

Fig. 3: Real and predicted arrival times in minutes of several stops from 11 Monday trajectories of driver A (A-C), and from 8 Friday trajectories of driver B (D).

4 Conclusion

In this paper, we have presented a Dynamic Neural Field approach to learning and memorizing ordinal and temporal properties of different driver routines. The learning is continuous starting with the first driving trajectory, happens implicitly without direct driver feedback, and can be scaled to different time periods (e.g., hours, days or weeks). The experiments with real GPS trajectory data validate the model architecture for two drivers sharing the car at different weekdays. The results confirm that the system correctly predicts in most cases the where and when of driver destinations. Prediction errors happen in cases of an exceptional, one-time deviation from the routine behavior. The activation-based learning mechanism implemented in the neural network also supports a fast adaption to consistent changes in the routines when predictions fail more than once or new destinations should be integrated in an existing routine. This contrast with many other neural network approaches (e.g., deep learning [5,25]) to learning sequential tasks which store the acquired information in synaptic weights. Weight-based learning typically requires large training sets and can be computationally very heavy. Moreover, personalized systems needs to be continuously updated and improved using cues from driver interaction. Such lifelong learning from non-stationary data distributions remains a long-standing challenge for machine learning and neural network models [17].

The integration of the long term memory with other factors such as traffic conditions in GPS-based navigation systems could be used to smarter route selection/recommendation without requiring input from the driver. Furthermore, the driver could benefit from advanced warnings about the possibility of arriving late at the next destination.

In future work will test the proposed model architecture in more complex driving scenarios, in the scope of the joint project "Easy Ride:Experience is everything" which we have with the car industry.

References

1. Amari, S.: Dynamics of pattern formation in lateral-inhibition type neural fields. Biological Cybernetics **27**(2), 77–87 (1977). https://doi.org/10.1007/BF00337259
2. Boukhechba, M., Bouzouane, A., Gaboury, S., Gouin-Vallerand, C., Giroux, S., Bouchard, B.: Prediction of next destinations from irregular patterns. Journal of Ambient Intelligence and Humanized Computing **9**(5), 1345–1357 (2018)
3. Eagle, N., Pentland, A.S.: Eigenbehaviors: Identifying structure in routine. Behavioral Ecology and Sociobiology **63**(7), 1057–1066 (2009)
4. Erlhagen, W., Bicho, E.: The dynamic neural field approach to cognitive robotics. Journal of Neural Engineering **3**, 36–54 (2006). https://doi.org/10.1088/1741-2560/3/3/R02
5. Fernandes, C., Ferreira, F., Erlhagen, W., Monteiro, S., Bicho, E.: A deep learning approach for intelligent cockpits: Learning drivers routines. In: International Conference on Intelligent Data Engineering and Automated Learning. pp. 173–183. Springer (2020)

6. Ferreira, F., Erlhagen, W., Bicho, E.: Multi-bump solutions in a neural field model with external inputs. Physica D: Nonlinear Phenomena **326**, 32–51 (2016). https://doi.org/10.1016/j.physd.2016.01.009
7. Ferreira, F., Erlhagen, W., Sousa, E., Louro, L., Bicho, E.: Learning a musical sequence by observation: A robotics implementation of a dynamic neural field model. In: 4th International Conference on Development and Learning and on Epigenetic Robotics. pp. 157–162. IEEE (2014)
8. Ferreira, F., Wojtak, W., Erlhagen, W., Vicente, P., Patel, A., Monteiro, S., Bicho, E.: A dynamic neural model for endowing intelligent cars with the ability to learn driver routines: where to go, when to arrive and how long to stay there. In: Towards Cognitive Vehicles Workshop (TCV2019), IROS2019. pp. 15–18 (2019)
9. Ferreira, F., Wojtak, W., Fernandes, C., Guimarães, P., Monteiro, S., Bicho, E., Erlhagen, W.: Dynamic identification of stop locations from GPS trajectories based on their temporal and spatial characteristics (2021), under review
10. Ferreira, F., Wojtak, W., Sousa, E., Louro, L., Bicho, E., Erlhagen, W.: Rapid learning of complex sequences with time constraints: A dynamic neural field model. IEEE Transactions on Cognitive and Developmental Systems (2020). https://doi.org/10.1109/TCDS.2020.2991789
11. Gonzalez, M.C., Hidalgo, C.A., Barabasi, A.L.: Understanding individual human mobility patterns. Nature **453**(7196), 779–782 (2008)
12. Hasenjäger, M., Heckmann, M., Wersing, H.: A survey of personalization for advanced driver assistance systems. IEEE Transactions on Intelligent Vehicles **5**(2), 335–344 (2019)
13. Hasenjäger, M., Wersing, H.: Personalization in advanced driver assistance systems and autonomous vehicles: A review. In: 2017 ieee 20th international conference on intelligent transportation systems (itsc). pp. 1–7. IEEE (2017)
14. Jiang, S., Ferreira, J., González, M.C.: Clustering daily patterns of human activities in the city. Data Mining and Knowledge Discovery **25**(3), 478–510 (2012)
15. Kun, A.L., et al.: Human-machine interaction for vehicles: Review and outlook. Foundations and Trends® in Human–Computer Interaction **11**(4), 201–293 (2018)
16. Laing, C.R., Troy, W.C., Gutkin, B., Ermentrout, G.B.: Multiple bumps in a neuronal model of working memory. SIAM Journal on Applied Mathematics **63**(1), 62–97 (2002). https://doi.org/10.1137/S0036139901389495
17. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S.: Continual lifelong learning with neural networks: A review. Neural Networks **113**, 54–71 (2019)
18. Rinzivillo, S., Gabrielli, L., Nanni, M., Pappalardo, L., Pedreschi, D., Giannotti, F.: The purpose of motion: Learning activities from individual mobility networks. In: 2014 International Conference on Data Science and Advanced Analytics (DSAA). pp. 312–318. IEEE (2014)
19. Sandamirskaya, Y., Zibner, S.K., Schneegans, S., Schöner, G.: Using dynamic field theory to extend the embodiment stance toward higher cognition. New Ideas in Psychology **31**(3), 322–339 (2013)
20. Schöner, G.: Dynamical systems approaches to cognition. Cambridge handbook of computational cognitive modeling pp. 101–126 (2008)
21. Simmons, R., Browning, B., Zhang, Y., Sadekar, V.: Learning to predict driver route and destination intent. In: 2006 IEEE Intelligent Transportation Systems Conference. pp. 127–132. IEEE (2006)
22. Sinnott, R.W.: Virtues of the Haversine. S&T **68**(2), 158 (1984)
23. Song, C., Koren, T., Wang, P., Barabási, A.L.: Modelling the scaling properties of human mobility. Nature Physics **6**(10), 818–823 (2010)

24. Song, C., Qu, Z., Blumm, N., Barabási, A.L.: Limits of predictability in human mobility. Science **327**(5968), 1018–1021 (2010)
25. Xu, J., Zhao, J., Zhou, R., Liu, C., Zhao, P., Zhao, L.: Predicting destinations by a deep learning based approach. IEEE Transactions on Knowledge and Data Engineering **33**(2), 651–666 (2021)
26. Yi, D., Su, J., Hu, L., Liu, C., Quddus, M., Dianati, M., Chen, W.H.: Implicit personalization in driving assistance: State-of-the-art and open issues. IEEE Transactions on Intelligent Vehicles **5**(3), 397–413 (2019)