CDP7 Results SGA2 (D4.4.2 – SGA2)

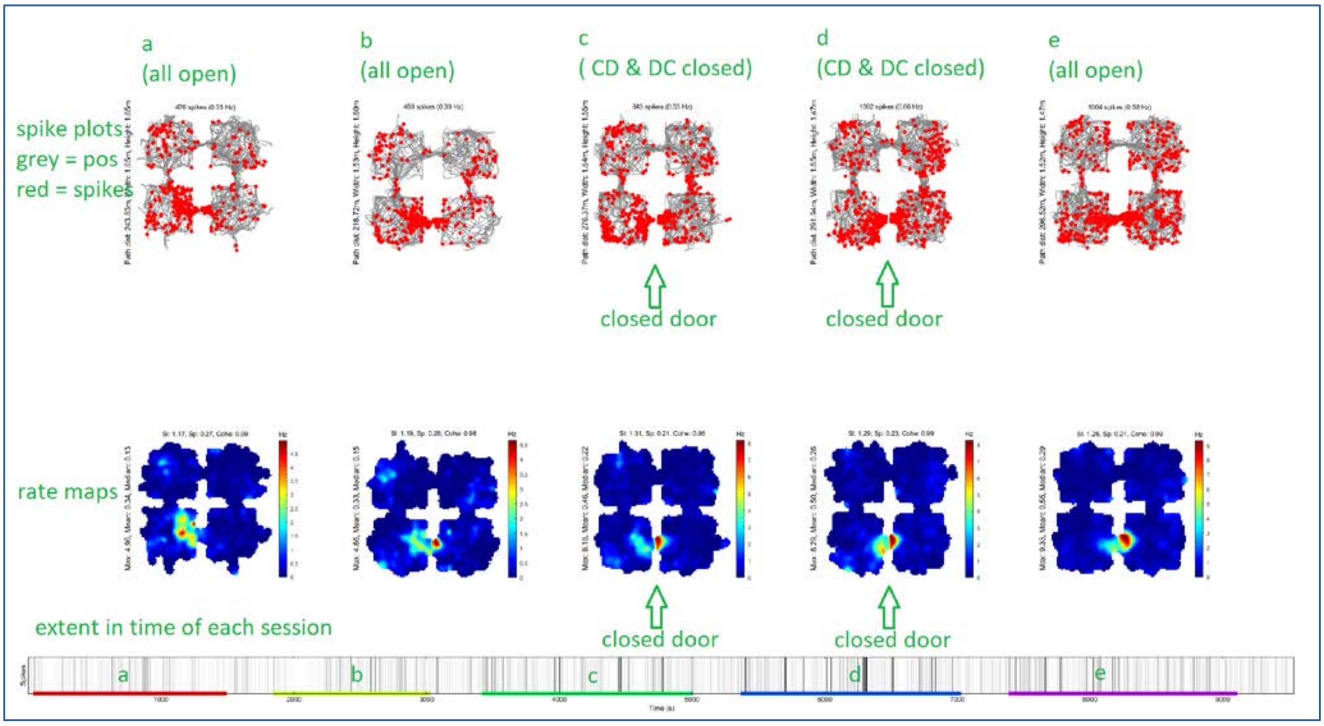

Figure 1: Example hippocampal (place) cell recorded while a rat navigates in a "four-room" scenario.

The Figure exemplifies the data that will help address some of the fundamental questions of CDP7, e.g., How does the hippocampus form spatial maps of (hierarchically) structured environments? How do these change during learning and when task structure changes (e.g. when a previously available passage is closed)? See for details Section 3.

Project Number:	785907	Project Title:	Human Brain Project SGA2

Document Title:	CDP7 Results SGA2
Document Filename:	D4.4.2 (D24.2 D122) SGA2 M24 ACCEPTED 201005.docx
Deliverable Number:	SGA2 D4.4.2 (D24.2, D112)
Deliverable Type:	Report
Work Packages:	Tasks T2.2.6, T2.2.7, T3.3.6, T4.3.4, T4.4.6, T4.5.4, T10.2.3
Key Result(s):	KRc7.1 / KRc7.2 / KRc7.3
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA2 M24 / 31 Mar 2020
Actual Delivery Date:	SGA2 M26 / 15 May 2020; resubmitted 31 Aug 2020; accepted 5 Oct 2020
Author(s):	Giovanni PEZZULO, CNR (P12); Christopher SUMMERFIELD, UOXF (P59); Hugo SPIERS, ULC (P82), Cyriel PENNARTZ, UvA (P98)
Compiled by:	Giovanni PEZZULO, CNR (P12)
Contributor(s):	Giovanni PEZZULO, CNR (P12); Chapters 1,2,5,6 Christopher SUMMERFIELD, UOXF (P59); Chapters 1,2,4,6 Hugo SPIERS, UCL (P82); Chapters 1,2,3,6 Cyriel PENNARTZ, UvA (P98); Chapters 1,2,6
SciTechCoord Review:	Martin TELEFONT, EPFL (P1)
Editorial Review:	Annemieke MICHELS, EPFL (P1)
Description in GA:	For consistent presentation of HBP results, SGA2 M24 Deliverables describing the accomplishments of an entire SP, WP or CDP have been prepared according to a standard template, which focuses on Key Results and the outputs that contribute to them. Project management elements such as Milestones and Risks will be covered, as per normal practice, in the SGA2 Project Periodic Report.
Abstract:	Humans and other animals are able to form sophisticated action plans; for example, planning routes to specific goal locations in spatial navigation. CDP7 studies the neuro-computational mechanisms of planning during spatial navigation in humans and rodents (as both species have excellent spatial navigation skills). In particular, it studies hierarchical planning - or how we subdivide our plans into "big chunks" or sub-goals (e.g. how to go from quarter to quarter in a city), rather than just consider step-by-step actions (e.g. specific sequences of left- and right-turns). To achieve these objectives, we designed spatial navigation tasks for humans and rodents, consisting in navigating and foraging in a hierarchically structured environment, composed of four rooms interconnected by corridors. Furthermore, we developed novel computational and robotic models that mimic navigation and planning abilities of humans and rodents. Our empirical results and computational models described within this document contribute to shed light into the mechanisms that we use to navigate and plan in structured environments.
Keywords:	Spatial navigation, hierarchical planning, hippocampus, medial temporal lobe, computational modelling

Target Users/Readers:	Computational neuroscience community, Consortium members, neuroimaging community, neuroscientific community, students

Table of Contents

Table of Tables

Table of Figures

Date	Change Requested / Change Made / Other Action
15 May 2020	Deliverable submitted to EC
30 Jul 2020	Resubmission with specified changes requested in Review Report Main changes requested: • Change 1: The deliverable includes in preamble a strange triple disclaimer that must be clarified before publication.
28 Aug 2020	Revised draft sent by SP/CDP to PCO. Main changes made, with indication where each change was made: • Change 1: The disclaimer has been removed (before table of contents) • Change 2: All headers and corresponding table titles were changed into more descriptive ones, for the sake of the reader (see all headers)
31 Aug 2020	Revised version resubmitted to EC by PCO via SyGMa

1. Overview

Humans and other animals are able to form sophisticated action plans; for example, planning routes to specific goal locations in spatial navigation. CDP7 studies the neuro-computational mechanisms of planning during spatial navigation in humans and rodents (as both species have excellent spatial navigation skills). In particular, it studies hierarchical planning - or how we subdivide our plans into "big chunks" or sub-goals (e.g. how to go from quarter to quarter in a city), rather than just consider step-by-step actions (e.g. specific sequences of left- and right-turns).

We designed spatial navigation tasks that both species (humans and rodents) can perform, and which consist in finding rewards in a hierarchically structured environment: 4 rooms interconnected by corridors (the hierarchical structure is evident if one considers that specific places can be considered either individually, or at a higher level of abstraction, as part of a room).

Furthermore, we developed novel computational models of planning during navigation, which we used to aid the analysis of human and rodent data, and used to control a robot that operates in the same "4-rooms" scenario as the humans and the rodents. The comparison of behaviour between living organisms (humans and rodents) and computational/robotic models is helpful in advancing our understanding of the neuro-computational principles of planning in complex environments.

Our novel experiments in the 4-rooms scenario strongly suggest that humans exploit hierarchical representations of space to generalise efficiently across tasks; and shed light on both the computational principles (e.g. inference over structured task representations) and the neuronal underpinnings (e.g. neural dynamics in medial temporal lobe and prefrontal cortex) supporting navigational planning.

In sum, the research conducted within CDP7 is advancing our understanding of the mechanisms of a key prospective ability - planning - that is still poorly understood at neural and computational levels.

2. Introduction

We have an incomplete understanding of how the brain of humans and other animals supports prospective and future-oriented forms of cognition. For example, we do not know what behavioural strategies we use to address large combinatorial planning problems that defy exhaustive search; and what neuro-computational mechanisms support such strategies.

CDP7 addresses the ways humans and rodents plan during spatial navigation, by using an innovative approach that integrates knowledge from these two species, as well as from computational and robotic models. It starts from the hypothesis that during navigation, humans and rodents plan using a hierarchical generative model of the task and its relevant states. To validate the hypothesis, CDP7 tested humans, rodents and computational/robotic models in a (largely) common spatial navigation scenario: a "4-room"-environment, which has been long considered in AI studies of hierarchical planning.

We performed a series of experiments that address both the behavioural signatures of planning during navigation and their neuronal underpinnings, using behavioural tasks, fMRI recordings (in humans) and single-cell hippocampal recordings (in rats). In parallel, we develop computational and robotic models of hierarchical planning during spatial navigation that both provide a quantitative characterisation of the project hypothesis and support model-based analysis of the experimental results.

The rest of the document is centred on the description of the Key Results and Outputs of CDP7. Section 3 describes the results of the rodent studies in the 4-room scenario. Two test conditions, when compared to a baseline condition, allow us to assess the effect of a change in environmental hierarchy and connectivity on hippocampal place cells. While the analysis of this dataset is still ongoing, preliminary results suggest that these changes do not have a major effect on the place cell population.

Section 4 describes the results of the human studies. In particular, the analysis of human behaviour indicates that humans exhibit strong forms of generalisation over hierarchical spatial structure. This strongly suggests that they are representing the 4-room environments hierarchically over multiple spatial scales, as predicted. This behavioural finding allows us to investigate the nature of the hierarchical representation in BOLD signals recorded whilst participants perform the task, with a focus on the medial temporal lobe and medial prefrontal cortex.

Section 5 describes the results of the computational modelling studies. In particular, Section 5.1.2 reports a novel computational model that characterises formally the neuronal circuit formed by the hippocampus and the ventral striatum; and explains how it supports goal-directed actions (from a probabilistic planning-as-inference perspective). Furthermore, it reports a novel computational model of the human study in the 4-room scenario described in Section 4. Section 5.1.3 reports the development and validation of two novel computational methodologies to support the analysis of spatial navigation data, at both behavioural and neural levels. Finally, Section 5.1.4 reports the development and validation of a robotic implementation of the above computational model of hippocampus and ventral striatum, within the Neurorobotic Platform (NRP) of SP10.

The above findings are important from a scientific perspective, as they contribute to shedding light into one of our most advanced cognitive abilities - planning - which is still poorly understood. Furthermore, the above findings are important from a technological perspective, given the pressing necessity to develop artificial systems that do more than react to the current stimuli, and are able to think about the future and the consequences of their actions (e.g. self-driving cars; robots that operate in rich social contexts).

3. Key Results KRc7.1 Hierarchical navigational planning - rodent neurophysiology

3.1 Outputs

3.1.1 Overview of Outputs

3.1.1.1 List of Outputs contributing to those models

- Output 1: Design of a 4-room environment and of a working task protocol (C3050)
- Output 2: Behavioural and electrophysiological dataset from hippocampal place cells in rats performing a "4-room" task (C3050)

3.1.1.2 How Outputs relate to each other and the models (link to the Key Results)

Output 1 is a prerequisite for Output 2.

3.1.2 Design of a 4-room environment and of a working task protocol

We designed a protocol to test the influence on rodent behaviour and hippocampal place cells of i) a change in the connectivity between rooms of a multi-compartment environment and ii) a change in the hierarchical organisation of these rooms. We also built and tested the actual maze in which this protocol would be run (Figure 2). In particular, we used doors that can be pushed open by the rats, which as far as we know have never been used before (Figure 3). These had to be carefully designed to make sure they were not too easy or too difficult to push, and that opening them did not impact the collection of neural signals. A detailed description of the design of the environment as well as of the paradigm will be included in a future research article.

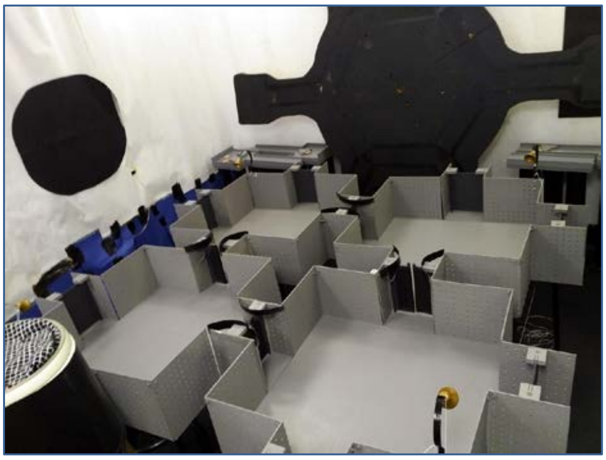

Figure 2: The 4-room maze

Figure 3: Example of rat pushing through door (source: Anyi LIU)

Figure 4: An implanted rat with a head-cap protecting the recording implant

For the protocol, we designed tests on the effect of a change of environmental connectivity / hierarchy on neural activity. Two different conditions are tested and the full protocol is presented in Table 1. Sessions are separated by breaks during which the rat is removed from the environment and placed on an elevated platform to rest with access to drinking water. For each session, the same sequence of rewarded boxes is used to try and make the behaviour as similar as possible between sessions.

Table 1: Protocol summary (test sessions highlighted with dashed lines)

Name	S1	S2	S3	S4	S5
“closed door”	All doors open	All doors open	1 door closed	1 door closed	All doors open
“one-way”	All doors open	All doors open	All doors open one-way	All doors open one-way	All doors open

By comparing the place cells’ activity between the “test” sessions (S3 and S4) and the “baseline” sessions (S1, S2 and S3) we can determine whether these changes are encoded in the activity of place cells. In the “closed-door” test, we expect local remapping, i.e. a change in the activity of cells which fire close to the place of the manipulation. In the “one-way” test, we expect either global remapping: most cells changing their place field (location of their firing) position, or general rate remapping: most cells changing their firing rate.

Table 2: Design of a 4-room environment and of a working task protocol Links

Component	Link to	URL
C3050	Data Repository	Data will be made available by the end of June 2020
	Technical Documentation	Documentation will be made available (as a preprint) by the end of June 2020
	User Documentation	Documentation will be made available (as a preprint) by the end of June 2020

3.1.3 Behavioural and electrophysiological dataset from hippocampal place cells in rats performing a "4-room" task

We trained 4 rats in the above-mentioned 4-room task and recorded electrophysiological signals from hippocampal pyramidal neurons ("Place cells"). Table 1 shows a preliminary count of place cells recorded in each condition. The recordings are ongoing with 3 more implanted rats to be recorded from, in order to collect enough data to perform robust analyses.

Table 3: Preliminary number of unique place cells recorded in each testing condition

Rat	"closed-door"	"one-way"
35	8	6
37	10	7
38	51	43
Total so far	**69**	**56**

Our first observations are that 1) rats know about the connectivity of the environment and use this information to navigate, as is shown by a decrease in the number of attempts to push the closed doors with time, and 2) place cells remain mostly spatially stable when the connectivity of the environment changes. This doesn't preclude subtler activity changes, induced by the connectivity / hierarchical change, that could be demonstrated by a more in-depth analysis of our dataset.

Two example place cells are illustrated in Figure 5. (**A**) A place cell recorded in a "closed-door" manipulation (sessions 3 and 4, top door is closed). Top: path of the rat (in grey) and cell activity (in red) in each session. Bottom: "rate maps" of the cell's activity in each session, i.e. firing rate at each location. Note how the spatial activity of this cell remains stable across sessions and even though a small activity change in activity seems to appear in the first test session (S3) this had actually already appeared in S2. (**B**) A place cell (from a different rat) recorded in a "one-way" manipulation (in S3 and S4, all 4 doors were closed anticlockwise, only allowing clockwise movements between rooms). Note how the spatial activity of this cell is also conserved across sessions and how, once again, a small activity change might be detected on the spike map in S3 but this is already visible in S2 so does not appear linked to our test manipulation.

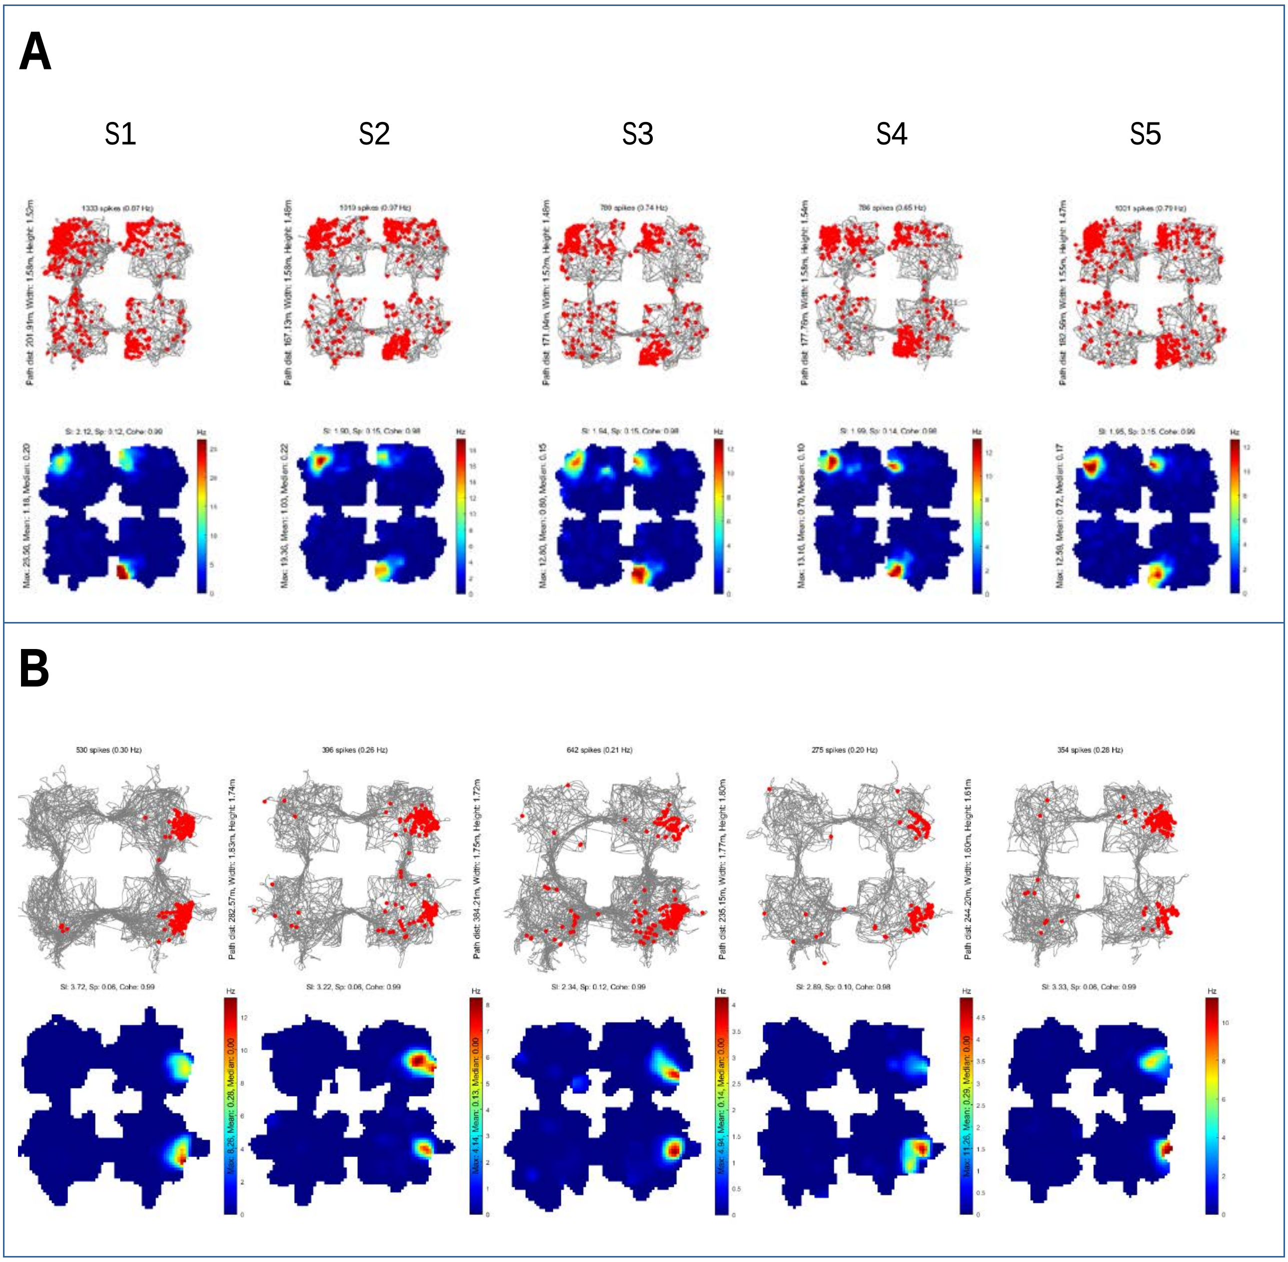

Figure 5: Example activity of place cells in the task, which show duplicated place fields in different rooms.

(A-B): Two example place cells in our dataset, which have multiple place fields. See Section 3.1.3 for details.

3.2 Validation and Impact

3.2.1 Actual and Potential Use of Output(s)

The designed protocol and collected data set are predicted to have an impact on both the spatial navigation scientific field as well as the computational modelling field. A critical manipulation was including conditions with open or closed doors, which permits assessing whether place cells remap as a function of changed transitions. In particular, the experiment permits comparing the predictions of different theories of hippocampal coding, including successor representation (SR), Boundary Vector Cell (BVC), and hierarchical coding (HC) models. For example, if CA1 place cells follow the SR model predictions, they should locally remap around a door with a changed connectivity; alternatively, if place cells are driven by BVCs, closing a door without changing its geometry should not affect place cells' firing significantly.

3.2.2 Publications

We expect to produce a biorXiv publication with this dataset by the end of June 2020.

4. Key Result KRc7.2 Hierarchical navigational planning - human neuroimaging

4.1 Outputs

4.1.1 Overview of Outputs

4.1.1.1 List of Outputs contributing to this KR

- Output 1: Design of a video game environment that involves hierarchical spatial planning, and the collection of human behavioural data (C3040)
- Output 2: Recording of fMRI data from a cohort of human participants during hierarchical spatial planning (C3040)
- Output 3: Behavioural data for navigation in a hierarchical maze (C3041).

4.1.1.2 How Outputs relate to each other and the Key Result

Output 1 is a prerequisite for Output 2. Output 3 is stand-alone.

4.1.2 Design of a video game environment that involves hierarchical spatial planning, and the collection of human behavioural data

We designed a 3D video game environment in which participants learned to forage for 2 successive rewards in a 4-rooms environment. A link to the task can be found here[1] and a representative image is shown below (Figure 6A).

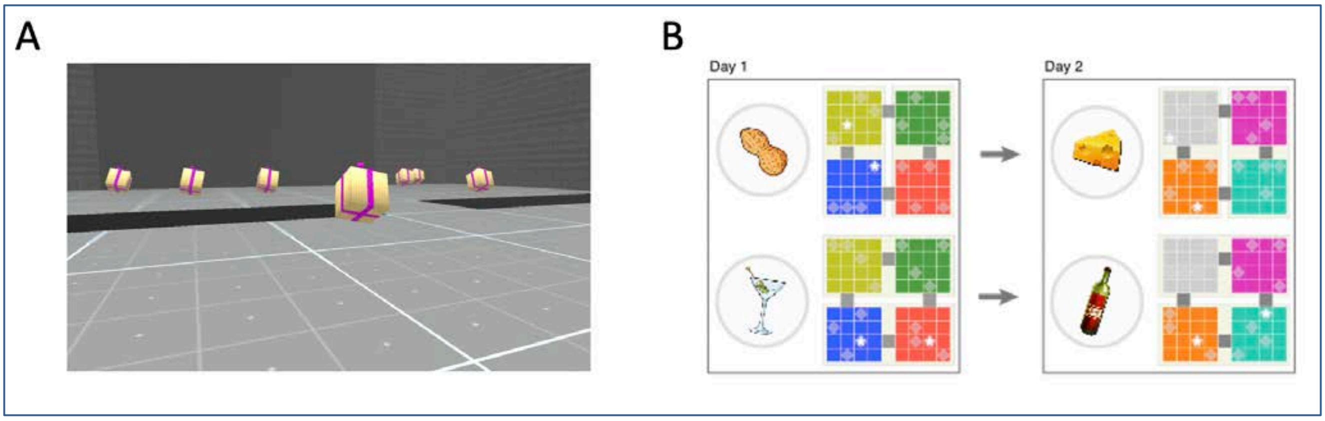

Figure 6: 3D video game environment for the human spatial navigation experiment.

The task required participants to move through the environment and open boxes, 2 of which contained a rewarding stimulus, that could be converted to a real financial incentive, paid as a bonus

[1] http://185.47.61.11/sandbox/tasks/hannahs/martinitask/dataset_1/builds/peanuts_martinis/

at the end of the experiment. Rewards had to be collected before a timeout. The relationship between the two rewards was signalled by a contextual cue that appeared at the start of the block: in one context, the rewards were in parallel locations along the vertical axis, and in another they were in parallel locations along the horizontal axis (Figure 6B). Participants then performed a new task the next day with differing room identifiers (floor colouring) but the same relational pattern. We constructed the task so that their actions after obtaining the first reward would betray whether they had learned and generalised the hierarchical spatial structure.

We collected human data from over 250 participants on this multi-session task. The data are made available via the Open Science Foundation website[2]. They consist of the following:

==
Experiment A part 1 (N=140): day1 blocked training
Experiment A part 2 (N=79): day2 blocked training on different reward types
==
Experiment B part 1 (N=24): day1 intermingled trial sequence training
Experiment B part 2 (N=16): day2 blocked training on different reward types

We aim to consolidate these data with neural measurements before publication. However, we already included some example data in Figure 7.

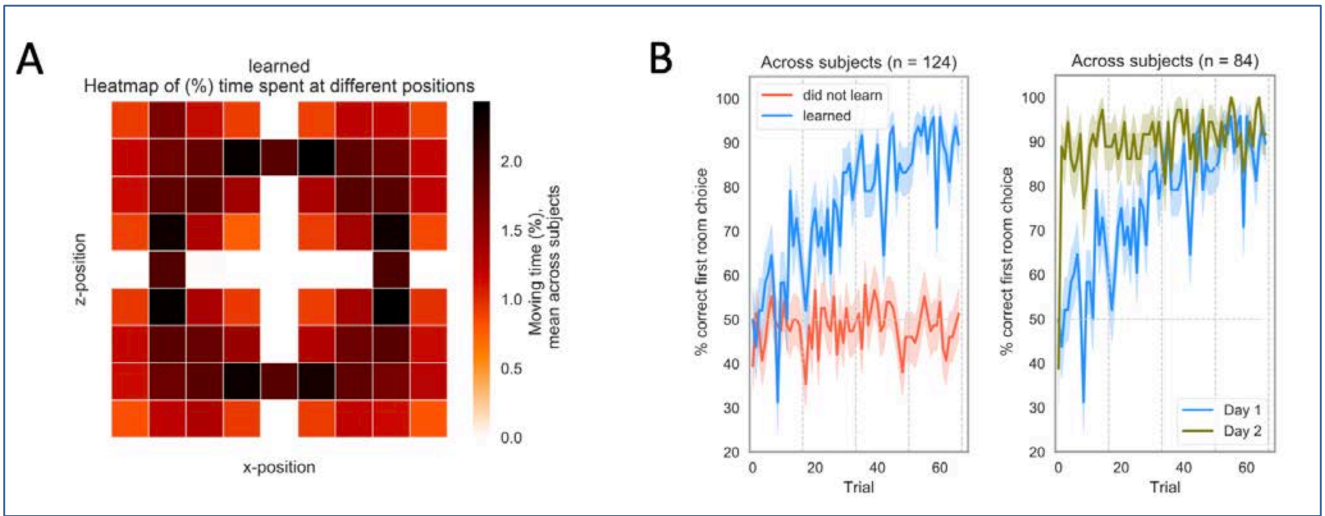

Figure 7: Example data from the human spatial navigation experiment.

(A) Heatmap of percent time spent by participants at different positions. (B) Percent of correct first-room choices across subjects.

Table 4: Design of a video game environment that involves hierarchical spatial planning, and the collection of human behavioural data - Links

Component	Link to	URL
C3040	Data Repository	Data are available via the Open Science Foundation website (https://osf.io/x6tge/) and EBRAINS https://wiki.ebrains.eu/bin/view/Collabs/sp2-collab/
	Technical Documentation	EBRAINS https://wiki.ebrains.eu/bin/view/Collabs/sp2-collab/
	User Documentation	EBRAINS https://wiki.ebrains.eu/bin/view/Collabs/sp2-collab/

[2] https://osf.io/x6tge/?view_only=2ff667d34a1d4f03b3575f09749e940a

4.1.3 Recording of fMRI data from a cohort of human participants during hierarchical spatial planning

We redesigned the 4-room task in a way that was optimised for the fMRI scanner. This involved changing the viewing perspective to a room-specific overhead view (Figure 6B) and introducing a "controller switch" whereby participants sometimes actively harvested rewards (as during behavioural training) and sometimes were moved through the maze towards rewards by an in-game AI controller. These careful design choices will allow us to conduct the analysis of spatial hierarchy in the hippocampus and other brain structures that constitute our work plan. Analyses of these data are ongoing. In Figure 8, we included an image of the fMRI task (Figure 8A) and of the medial temporal lobe ROI we are using (Figure 8B). We expect to have completed the analysis of the fMRI data before summer 2020.

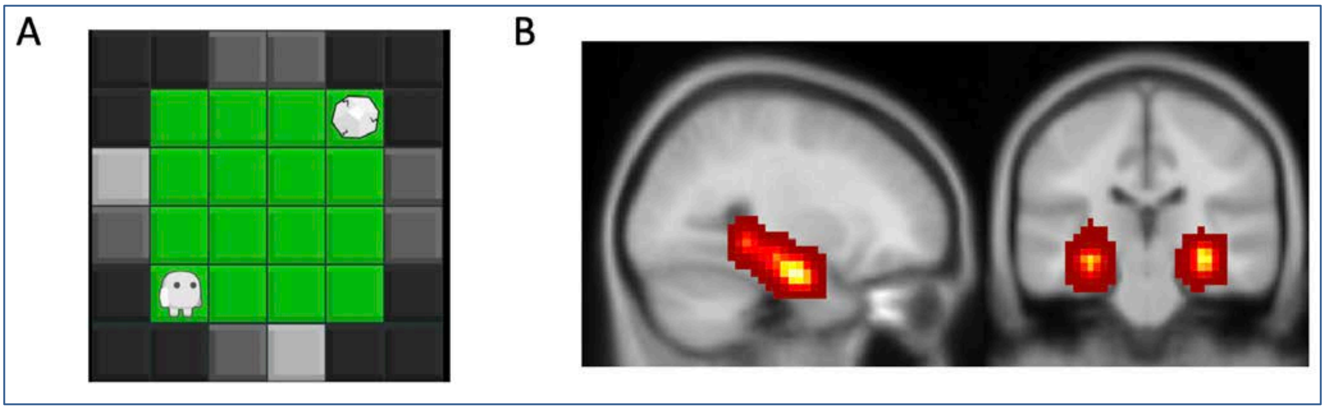

Figure 8: fMRI version of the human spatial navigation experiment.

(A) Screenshot of the fMRI task. (B) Medial temporal lobe ROI.

4.1.4 Behavioural data for navigation in a hierarchical maze

We designed and tested an experiment that involved learning from scratch to navigate in a hierarchical maze. The maze was symbolic, i.e. only signalled by object cues (Figure 9A). Different computational models make distinct predictions about how different training curricula should affect training. We tested these in a behavioural experiment involving >90 participants (Figure 9B). These data will be made publicly available shortly through the Open Science Foundation.

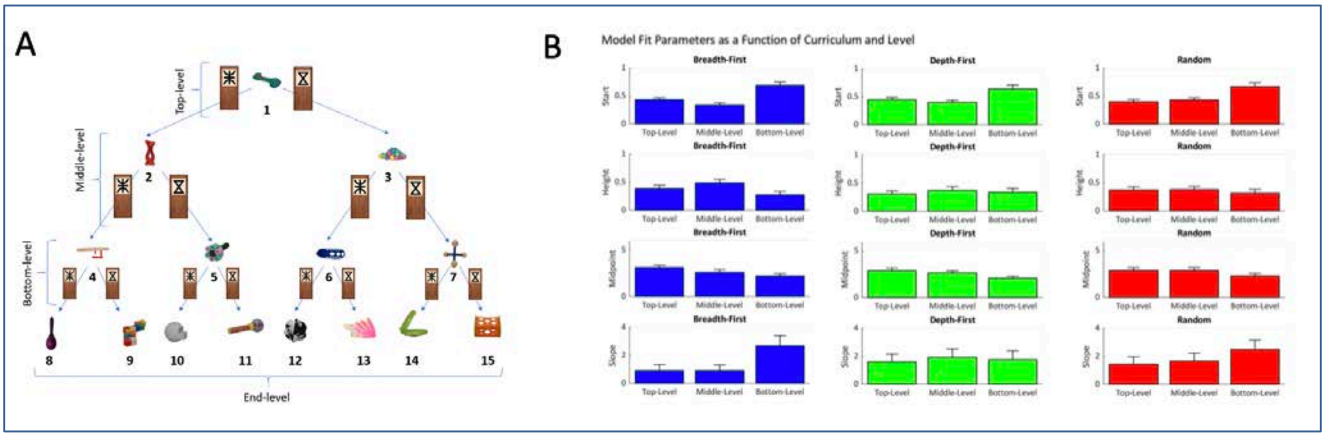

Figure 9: Human experiment on learning to navigate in a hierarchical maze.

(A) Object cues used in the study and their hierarchical structure. (B) Example (preliminary) results of the study.

Table 5: Behavioural data for navigation in a hierarchical maze - Links

Component	Link to	URL
C3041	Data Repository	Data will be made available through the Open Science Foundation and within the HBP platform by the end of June 2020
	Technical Documentation	Documentation will be made available within the HBP platform by the end of June 2020
	User Documentation	Documentation will be made available within the HBP platform by the end of June 2020

4.2 Validation and Impact

4.2.1 Actual and Potential Use of Output(s)

The designed protocol and collected data set will be of interest to communities in both systems and theoretical neuroscience, with a focus on navigation, planning and reasoning. In particular, the data permit validating alternative models of hierarchical coding and abstraction. One example hierarchical model used to characterise the experimental findings is provided in Section 5.1.2. It describes the abstraction process in terms of Bayesian nonparametric methods and successfully characterises the generalisation ability of human participants to the study described in Section 4.1.2.

4.2.2 Publications

We expect to publish these data in 2020.

5. Key Result KRc7.3 Hierarchical navigational planning - computational modelling

5.1 Outputs

5.1.1 Overview of Outputs

5.1.1.1 List of Outputs contributing to this KR

- **Output 1:** Model for look-ahead prediction and planning during spatial navigation (C3037)
- **Output 2:** Model-based analysis of neural data obtained during goal-directed spatial navigation (C3036)
- **Output 3:** Robotic model of hierarchical planning during goal-directed spatial navigation (C3035)

(Note the work reported here relates to Tasks T4.3.4, T4.4.6, T4.5.4, whose leader is CNR).

5.1.1.2 How Outputs relate to each other and the Key Result

Output 1 is a prerequisite for Output 3. Output 2 is stand-alone.

5.1.2 Model for look-ahead prediction and planning during spatial navigation

We developed a novel computational model of hippocampus-based spatial navigation, which tests the idea that the hippocampus and the ventral striatum jointly form a goal-directed controller – implemented in the model using Bayesian nonparametrics and planning-as-inference, see Figure 10. We showed that the computational model can address problems requiring look-ahead prediction and planning; and during learning, it develops internal codes having key characteristics of neurons in the hippocampus and the ventral striatum (as compared with single cell data from CDP7 members, PENNARTZ Lab). The model was fully described in Stoianov *et al*. (2018) -P1393 and the source code was released in open source format. In subsequent publications, we explored developments of the model that cover additional mechanisms beyond goal-directed spatial navigation, and namely behavioural habitisation (Maisto *et al.* (2019) - P1957) and exploration (Pezzulo and Nolfi (2019) - P1823).

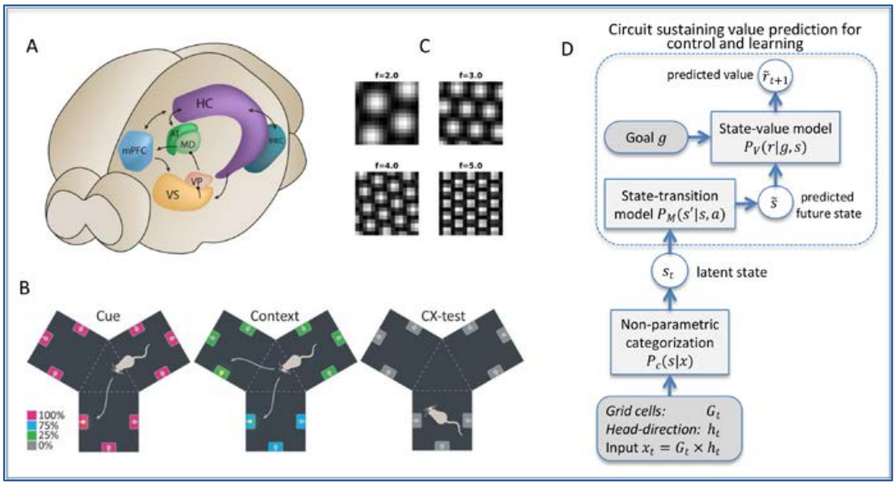

Figure 10: Schematic of the computational model of goal-directed navigation.

(A) Sketch of systems-level circuit investigated in the study. (B) Sketch of the behavioural paradigm. (C) Example grid cells used in the model. (D) Illustration of the model, see Stoianov *et al*. (2018) -P1393 for details.

Furthermore, we developed a novel computational model of the human spatial navigation task described in Section 4.1.2 (the 4-room scenario). The model leverages and extends the Bayesian nonparametric model of spatial navigation illustrated in Figure 10, with an additional mechanism that acquires hidden task rules (e.g. the fact that reward locations may be arranged horizontally or vertically). Figure 11 (A) shows the structure of the model, with three main components that develop latent internal codes based on contextual stimuli and learn their probabilistic dependencies - which afford probabilistic inference about task rules and reward location. Figure 11 (B) shows the behavioural results of the model, with an immediate transfer of learned task rules from day 1 to day 2, as observed experimentally in humans (see Output 1 of Section 4). Figure 11 (C) shows the analysis of latent codes (S, Z and Y) developed in the three components of the model, using measures of mutual information (MI) and conditional mutual information (CMI). This analysis illustrates that the three components of the model develop latent codes that are sensitive to different kinds of information - that have to be integrated to successfully solve the task. These analyses are useful to compare with information content in different brain areas, as revealed by fMRI. The computational model is still unpublished and the current draft is [online here][3]. We expect to publish a paper comparing hidden variables of the model (as emerged from the Bayesian nonparametric model) with neuronal codes in medial temporal lobe and prefrontal cortex (from the human fMRI experiment, see Section 4.1.3) before summer 2020.

[3] https://sites.google.com/site/giovannipezzulo/home/publications/filecabinet/CDP7-four_rooms_task%20.pdf

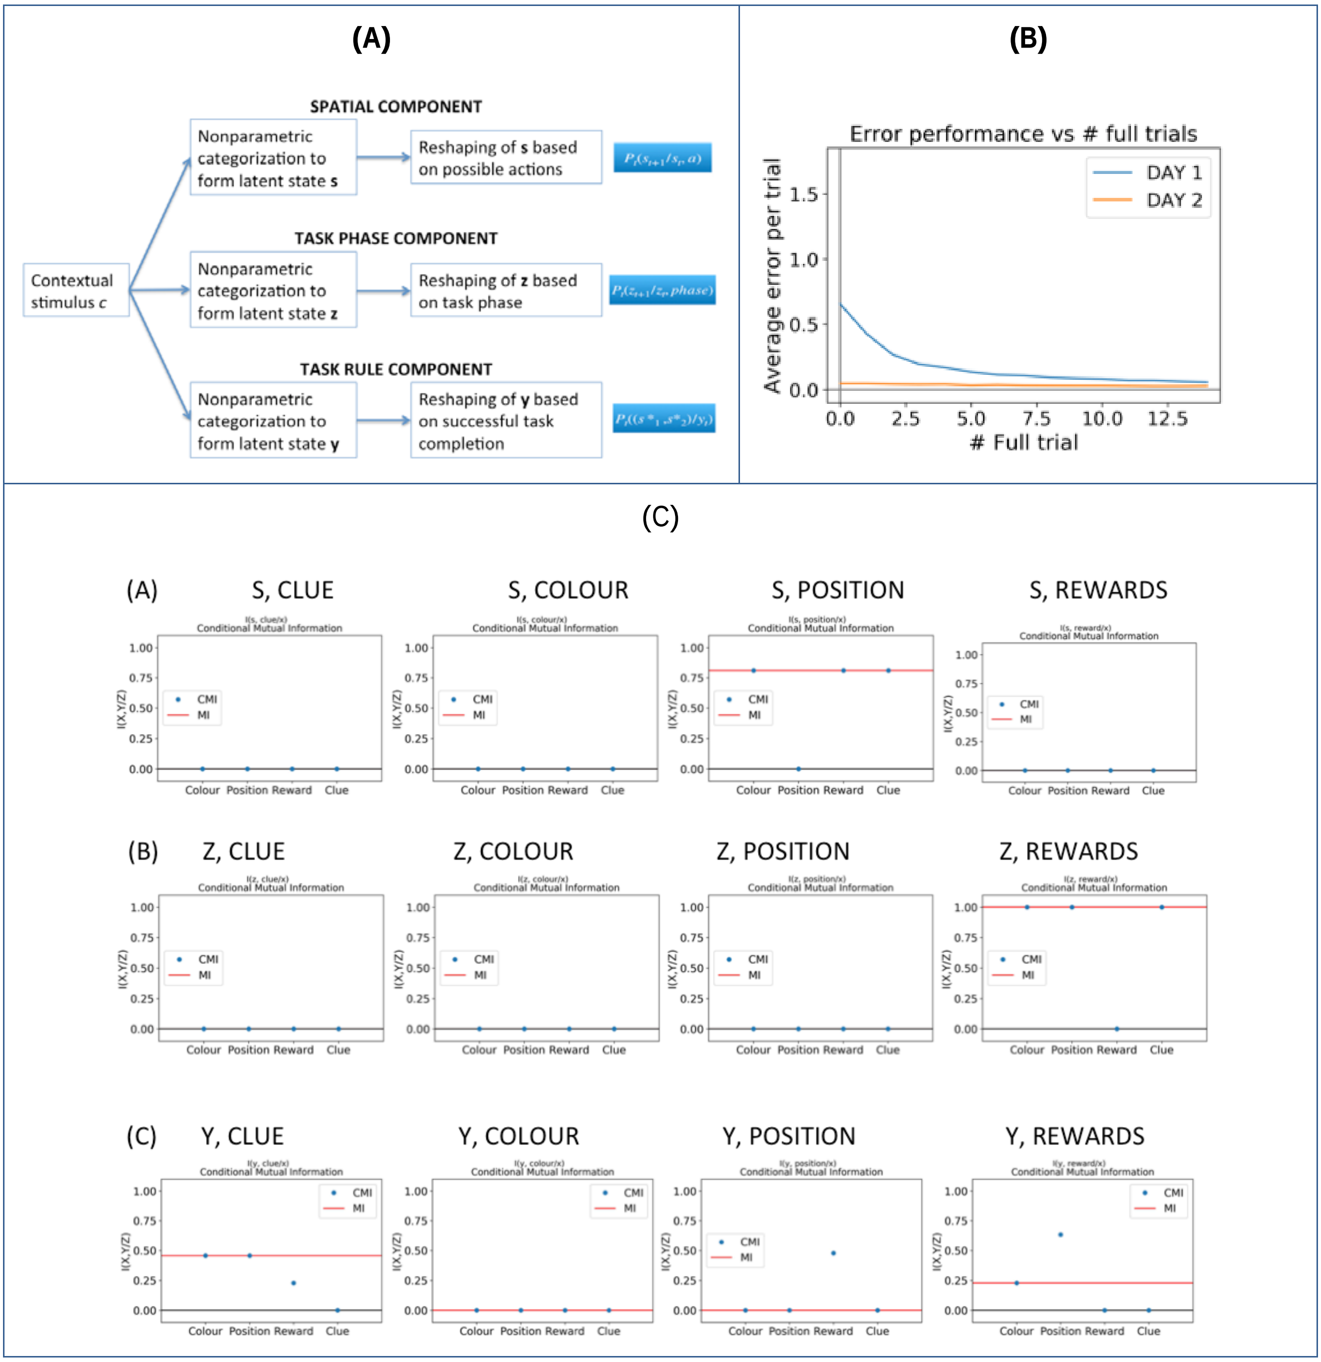

Figure 11: Schematic of computational model of human spatial navigation in 4-rooms tasks

Schematic of computational model of human spatial navigation in 4-rooms tasks, reported in Section 4.1.2. (A) Sketch of the main components of the model. (B) Performance of the model during days 1 (blue) and 2 (red). (C) Analysis of the information content of latent states developed by the computational model. See[3] for additional details.

Table 6: Model for look-ahead prediction and planning during spatial navigation - Links

Component	Link to	URL
C3037	Data Repository	https://kg.ebrains.eu/search/instances/Model/ed0554b569a59e3ccabe81eb592dfc42e216140e https://github.com/stoianov/MBRL/tree/1.0
	Technical Documentation	https://kg.ebrains.eu/search/instances/Model/ed0554b569a59e3ccabe81eb592dfc42e216140e
	User Documentation	https://kg.ebrains.eu/search/instances/Model/ed0554b569a59e3ccabe81eb592dfc42e216140e

5.1.3 *Model-based analysis of neural data obtained during goal-directed spatial navigation*

This Output includes two novel computational methodologies for the analysis of spatial navigation data collected within CDP7 and beyond, at behavioural and neural levels.

The former is a novel computational model that directly incorporates the main hypothesis of CDP7, by casting the hippocampal formation as a (hierarchically structured) generative model and hippocampal spontaneous dynamics as generative replays from the model (Figure 12). In addition to illustrating a novel theory of hippocampal function (and especially hippocampal replays), the computational model can be used to perform model-based analyses of the statistics of hippocampal codes (e.g. place cells) and of replays observed while rodents navigate in structured environments (e.g. the 4-room scenario) and when they rest afterwards. The computational model is reported in the preprint Stoianov *et al*. (2020) - P2331; while the results of the model-based analyses of neural data are preliminary and yet unpublished.

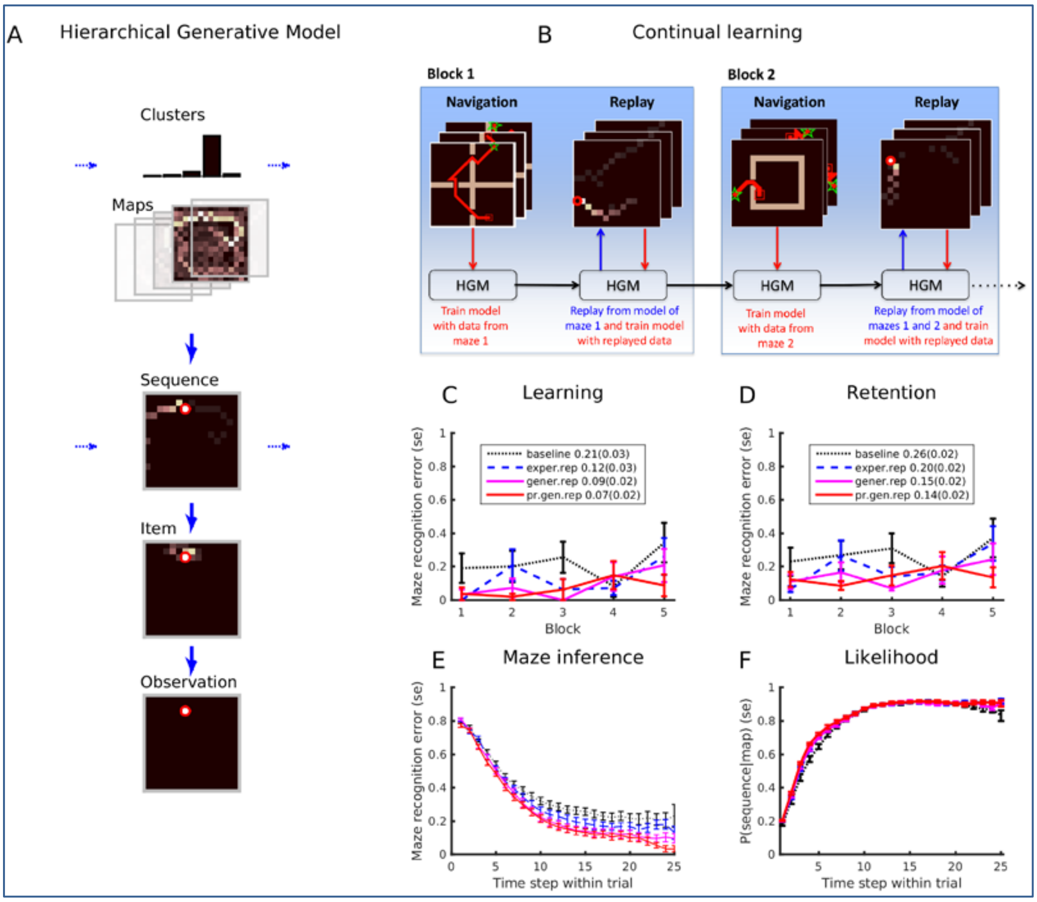

Figure 12: Schematic of the model of hippocampal formation as generative model.

(A) Structure of the model. (B) Set up of the continual learning experiment. (C-D) Performance of the model. See Stoianov et al. (2020) -P2331 for additional details.

The latter is a novel method to identify "motor primitives" in rodent spatial trajectories (i.e., elementary movement units that can be composed to reconstruct the animal's movements in space) using a machine learning technique: dictionary learning. We validated the "motor primitives" approach using an available dataset of rodent movements. Our results are reported in the preprint Donnarumma *et al*. (2020) - P2417 (*in validation process*). They show that the methodology permits the identification of structured behavioural patterns within rodent spatial trajectories (see Figure 13); the identification of specific characteristics of the maze (e.g. complexity) or the animals' state (e.g. stereotyped behaviour) from movements; and the prediction of place and grid displacement in novel mazes. Figure 13 illustrates some aspects of the model; and namely, examples of motor primitives (a); fictive trajectories generated from the motor primitives (b); predicted displacement of grid (c-d) and place cells (e-f) during spatial navigation. This methodology can be used by behavioural scientists and neuroscientists as an aid for behavioural and neural data analysis.

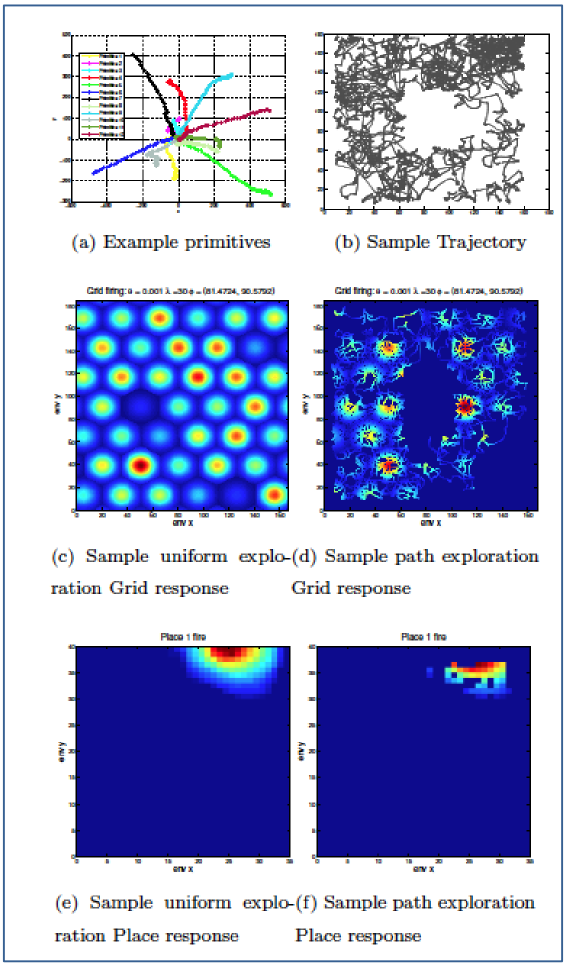

(a) Example primitives (b) Sample Trajectory

(c) Sample uniform exploration Grid response (d) Sample path exploration Grid response

(e) Sample uniform exploration Place response (f) Sample path exploration Place response

Figure 13: Simulated data from computational model of rodent spatial trajectories

(see Donnarumma *et al*. (2020) - P2417 *(in validation process)*).

Table 7: Model-based analysis of neural data obtained during goal-directed spatial navigation - Links

Component	Link to	URL
C3036	Data Repository	https://kg.ebrains.eu/search/instances/Model/5d2ed7e2545226fa316f75774fd020fd https://github.com/stoianov/HDGM/tree/V1
	Technical Documentation	https://kg.ebrains.eu/search/instances/Model/5d2ed7e2545226fa316f75774fd020fd
	User Documentation	https://kg.ebrains.eu/search/instances/Model/5d2ed7e2545226fa316f75774fd020fd

5.1.4 Robotic model of hierarchical planning during goal-directed spatial navigation

We embodied the computational model of goal-directed navigation (described as Output 3) in a Husky robot, simulated on the Neurorobotic Platform of SP10, in order to assess the scalability and robustness of the model and extend it to include more realistic (robotic) action-perception control loops. This work was conducted in collaboration with HBP partner Scuola Superiore Sant'Anna (SSSA);

further details are specified in SGA2 Deliverable D10.5.2 (D67.2 D48). The Husky robot was successfully tested in three different navigation scenarios, where it had to learn to navigate autonomously to goal locations (in red in Figure 14). The results of this work have not been published yet. However, they have been presented in a poster at the HBP Summit 2020 (M. Priorelli, M. Kirtay, I. P. Stoianov, G. Pezzulo, E. Falotico "The hippocampus-ventral striatum circuit model in the Neurorobotics Platform in navigation tasks") and in a conference paper accepted at the I-RIM conference https://i-rim.it/it/conferenza-i-rim/.

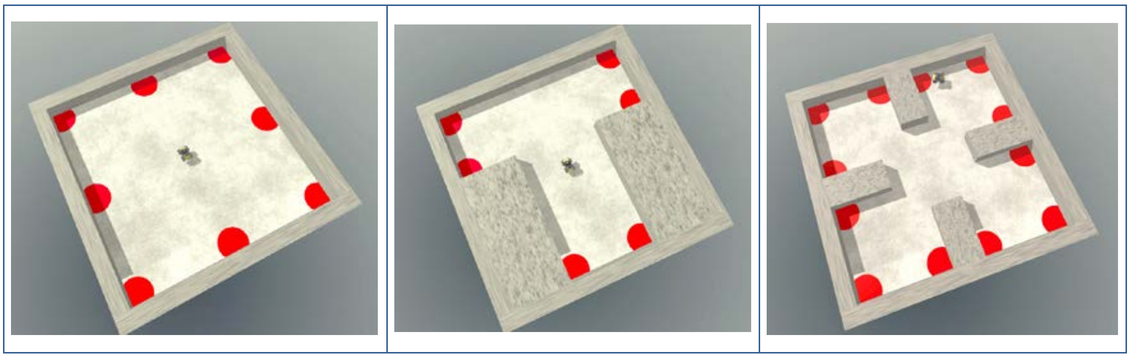

Figure 14: Three scenarios where the neurorobotic model was tested.

Table 8: Robotic model of hierarchical planning during goal-directed spatial navigation - Links

Component	Link to	URL
C3035	Data Repository	https://collab.humanbrainproject.eu/#/collab/78682/nav/532632
	Technical Documentation	https://collab.humanbrainproject.eu/#/collab/78682/nav/532632
	User Documentation	https://collab.humanbrainproject.eu/#/collab/78682/nav/532632

5.2 Validation and Impact

5.2.1 Actual and Potential Use of Output(s)

The impact of the work conducted under this KRc7.3 is mainly scientific: it advances our knowledge of the computational and neural mechanisms supporting goal-directed navigation and planning; it extends the toolbox of data analysis methods (to target especially neural codes); and it demonstrates robotic implementations of autonomous, goal-directed navigation.

5.2.2 Publications

- P1393 - Stoianov I., Pennartz, C., Lansink, C., Pezzulo, G. (2018) Model-based spatial navigation in the hippocampus - ventral striatum circuit: A computational analysis. *PLoS Computational Biology*, 14(9) e1006316

  Highlight for Output 1: This publication reports a novel computational model of spatial navigation, where hippocampal and ventral striatal dynamics implement look-ahead and planning, realised in collaboration with SP3 (PENNARTZ Lab). This publication validates **Output 1** with scientific peer review. Furthermore, the model described in this publication was implemented in the neurorobotic platform (**Output 3**).

- P2331 - Stoianov I., Maisto D., Pezzulo G., (2020) The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *BiorXiv preprint*

  Highlight for Output 2: This preprint publication introduces a hippocampal model that supports a dual use: first, it provides a novel theoretical explanation of cognitive map formation and hippocampal replays; second, it provides a novel computational method for the analysis of latent representations formed during spatial navigation and of hippocampal spontaneous activity (**Output 2**).

- P2417 *in validation process*- Donnarumma F., Prevete F., Maisto D., Fuscone A., van der Meer M., Kemere C., Pezzulo G. (2020) A framework to identify structured behavioral patterns within rodent spatial trajectories. *Biorxiv preprint*

  This preprint publication introduces a novel methodology to extract motor primitives from rodent spatial trajectories, to be used to analyse animal data at both the behavioural level (e.g., identification of path stereotypy and preferences) and the neural level (e.g. prediction of place and grid displacement in novel mazes). (**Output 2**)

- P1957 - Maisto D., Friston K., Pezzulo G. (2019) Caching Mechanisms for Habit Formation in Active Inference. *Neurocomputing* 359, 298-314

  This publication provides a novel computational perspective on how behavioural habits are formed after (over)training and how they complement goal-directed navigation mechanisms based on the hippocampus. It therefore elaborates and extends the results of **Output 1** beyond goal-directed spatial navigation, to also address behaviour routinisation.

- P1823 - Pezzulo, G., Nolfi, S. (2019) Making the Environment an Informative Place: A Conceptual Analysis of Epistemic Policies and Sensorimotor Coordination. *Entropy* 21 (350)

  This publication provides a computationally-guided analysis of epistemic and exploratory behaviour, of the kind involved in the formation of a task space or a (hippocampal) cognitive map. It therefore elaborates and extends the results of **Output 1** beyond goal-directed spatial navigation, to also address exploration.

- P2130 *in validation process*- Pezzulo G., Donnarumma F., Maisto D., Stoianov I. (2019) Planning at decision time and in the background during spatial navigation. *Current Opinion in Behavioral Science*, 29, 69-76

  This review provides a novel conceptual overview of the neural and computational mechanisms supporting spatial navigation and planning, with a particular focus on the hippocampus. It is intended to present the results of **Outputs 1, 2 and 3** to a wide scientific audience.

- P1959 - Maranesi M., Bruni S., Livi A., Donnarumma F., Pezzulo G., Bonini L. (2019) Differential neural dynamics underling pragmatic and semantic affordance processing in macaque ventral premotor cortex. *Scientific reports* 9 (1), 1-11

  This paper widens the scope of the computational data analysis of **Output 2**, by targeting the neural mechanisms for planning object-directed actions (in monkeys).

- P2006 - Gómez CM., Arjona A., Donnarumma F., Maisto D., Rodriguez Martinez E.I., Pezzulo G. (2019) Tracking the time course of Bayesian inference with Event Related Potentials: a study using the central cue Posner paradigm. *Frontiers in Psychology* 10:1424. doi: 10.3389/fpsyg.2019.01424

  This paper widens the scope of the computational data analysis of **Output 2**, by targeting the neural mechanisms that permit estimating hidden context and surprise, in humans.

6. Conclusion and Outlook

In CDP7, we studied navigational planning in humans and rodents by combining behavioural, neural, computational modelling and robotic approaches. The central questions we addressed were whether living organisms (humans or rodents) organise their planning hierarchically (e.g. plan in terms of sub-goals rather than in terms of step-by-step action sequences) and what kind of neuro-computational mechanisms are required for such hierarchical planning (e.g. forms of structure learning that permit developing useful task abstractions).

To address these challenges, we designed coherent human, rodent and robotic tasks in a common spatial navigation scenario – called the “4-room” scenario – which has been long considered an ideal scenario to study hierarchical plans in AI.

Our experiments in CDP7 tested whether (and how) rodents, humans and computational models/robots represent hierarchical task structure and use it for adaptive decision and generalisation to novel situations where (part of the) structure is preserved. The focus on two different species (humans and rodents) and computational/robotic models permits addressing different, complementary facets of this question.

Our rodent experiments in the 4-room scenario aim at elucidating the neuronal underpinnings supporting hierarchical spatial coding and planning, by testing the effect of a change in environmental hierarchy and connectivity on hippocampal place cells; see Section 3. While our analyses are still ongoing, the preliminary data suggest that the most immediate aspects of place cell coding (place fields) are not immediately influenced by changes of connectivity. Yet we have developed a suite of computational methods that permit testing sequential coding in the hippocampus, beyond place fields; and these novel analysis tools (that we are testing) hold the promise to unveil the mechanisms permitting rodents to flexibly reuse and generalise spatial knowledge.

Our human experiments in the 4-room scenario tested the ability of humans to learn hidden task regularities (e.g. that pairs of rewards could be found at predictable locations, arranged horizontally or vertically) and exploit them to generalise to novel situations, where some superficial features of the task (e.g. colour and reward identity) changed, but hidden regularities (e.g. horizontal or vertical arrangement) did not; see Section 4. Our results strongly suggest that we humans use hierarchical spatial structure to generalise rapidly to novel situations that have analogous (hidden) regularities. The behavioural results also paved the way to the development of an fMRI version of the 4-room scenario, which aims to shed light on the neuronal mechanisms supporting the representation and use of hierarchical task structure (analyses are ongoing).

Our computational modelling and robotic experiments complement the human and rodent experiments, by permitting to develop and test mechanistic hypotheses on hierarchical task representations and planning. These activities comprised several computational models that test spatial navigation from different perspectives; see Section 5. One model addresses the hippocampus – ventral striatum circuit during spatial navigation, and was successfully embodied in the Neurorobotic Platform (NRP), hence demonstrating its scalability and potential to explain sophisticated navigational skills. A second model directly simulates human data on spatial navigation in the 4-room task collected in CDP7, in order to investigate the putative computational mechanisms affording hierarchical structure learning and generalisation in humans. When tested in the 4-room scenario with a protocol analogous to the human experiment, this second model shows the same transfer learning ability as humans and develops latent codes that code for key spatial and structural property of the task. This model therefore suggests candidate mechanisms for human structure learning and generalisation, which we are testing by comparing and "aligning" model predictions and human data. A third model directly implements one of the key assumptions of CDP7 – that the hippocampus functions as a generative model, to learn hierarchical spatial codes for maps and sequences – and provides a novel perspective on internally generated hippocampal sequences (and replays).

Taken together, these rodent, human and computational studies offer a number of novel results and insights that greatly advance our understanding of planning during spatial navigation – and more broadly, how we learn task structure to generalise to novel tasks and think about the future. The collaborative work conducted during CDP7 has already produced several novel theoretical, empirical and computational advancements, as testified by the results shown in this report and the accompanying papers (published or to be released soon). Given the time constraints of CDP7 (which lasted only 2 years), some implications of our studies remain to be fully developed and some analyses are still ongoing; including most prominently the analyses of neural data – which will be consolidated and finalised within SGA3.

Another impact of CDP7 relates to scientific dissemination. Three PIs of CDP7 (Giovanni PEZZULO, Cyriel PENNARTZ, Chris SUMMERFIELD) together with another HBP partner (Lars MUCKLI) co-organised a 2-day workshop entitled "Predictive coding, inference and unsupervised learning" at the European Institute for Theoretical Neuroscience, 16-17 Jan 2019. Details on the workshop can be found here: https://eitnconf-160119.sciencesconf.org/