Project Number:	284941	Project Title:	Human Brain Project

Document Title:	Applications: progress report (RESUBMITTED)
Document Filename:	SP11 D11.4.2 RESUBMITTED FINAL.docx
Deliverable Number:	D11.4.2 (RESUBMITTED)
Deliverable Type:	Report
Work Package(s):	WP11.1, WP11.2, WP11.3
Dissemination Level:	PU = Public
Planned Delivery Date:	M12/30 September 2014
Actual Delivery Date:	SUBMITTED: M15/15 December 2014 RESUBMITTED: M22/29 July 2015

Authors:	Karlheinz MEIER UHEI (P45)	*Subproject Leader, WP11.3*
Compiling Editors:	Martina SCHMALHOLZ, UHEI (P45)	*Compiled Deliverable*
Contributors:	Greg FRANCIS, EPFL (P1)	*WP11.1, T11.1.1, T11.1.2*
Contributors:	Ferath KHERIF, CHUV (P23)	*WP11.2, T11.2.1*
Contributors:	Frank GOTTFRIED, SAP (P47)	*T11.3.1*
Contributors:	Christian HUYCK, MU (P91)	*T11.3.2*
Contributors:	Bernabe LINARES, CSIC (P82)	*T11.3.3*
Contributors:	Ulrich RÜCKERT, UNIBI (P102)	*T11.3.4*
Contributors:	Ryad BENOSMAN, UPMC (P107)	*T11.3.5*
Contributors:	Thomas NOWOTNY, UOS (P112)	*T11.3.6*
STO Review:	UHEI (P45): Martina SCHMALHOLZ	
Editorial Review:	EPFL (P1): Guy WILLIS, Lauren ORWIN	

Abstract:	This Deliverable describes the work and the achievements of SP 11 (applications) during Year 1. SP11 has progressed according to the original description of work. An additional effort beyond the DoW has been the integration of five new groups from the Competitive Call into the Future Computing Work Package. SP11 has a close connection to other SPs; the HBP Platforms in particular. This connection is visible in all activities throughout the first Year 1. Preparation of Platform applications will continue in all three scientific Work Packages in Year 2.
Keywords:	Applications, Future Neuroscience, Future Medicine, Future Computing, integration, cognitive neuroscience, clinical data, data mining, neuromorphic computing, NM-PC-1, NM-MC-1

Table of Contents

The HBP Review in January 2015 required the resubmission of D11.4.2 “Applications: Progress Report”, with additional progress descriptions for specific Tasks, and details of the resources used by these Tasks. SP11 recognises the importance of the points raised by the Review, and addresses them below.

Executive Summary

Subproject 11 (Applications) has the important role of preparing use cases for the HBP Platforms in the three Project research areas: Future Neuroscience, Future Medicine and Future Computing. It has to follow the process of Platform building, design the Applications accordingly and start implementing them towards the end of the Ramp-Up Phase. At the end of the Ramp-Up Phase, the Platforms will analyse the SP11 Use Cases and propose possible changes arising from the SP11 results.

SP11's main achievement is that all three of its scientific Work Packages have advanced well in defining their Applications. In particular, the Future Computing research area has successfully integrated the five new groups that were selected via the Competitive Call. As these groups had only seven months to accomplish the work scheduled for the first Project Year and had not been involved in the Proposal Writing Phase, this is a remarkable achievement.

Two out of seven Milestones were not reached in time for Month 12 and have been delayed by a few weeks. The reasons are clearly described by the groups responsible and will not have any detrimental effect on overall progress. In one case, the reason for the delay was the late receipt of funding (MS315). In another, the delay was due to the unexpectedly large number of network architectures that had to be evaluated (M209).

During the next six months, work will continue towards implementation. After Month 18, the process of commissioning the Platforms will begin and the Applications Work Packages will be involved in this as they implement their Use Cases.

1. Introduction

1.1 The Human Brain Project (HBP)

The Human Brain Project (HBP) is a major international scientific research project, involving over 100 academic and corporate entities in more than 20 countries. Funded by the European Commission (EC), the ten-year, EUR 1 billion Project was launched in 2013 with the goal "to build a completely new ICT infrastructure for neuroscience, and for brain-related research in medicine and computing, catalysing a global collaborative effort to understand the human brain and its diseases and ultimately to emulate its computational capabilities."

The fields of neuroscience, medicine and information technology each have important roles to play in addressing this challenge, but the knowledge and data that each is generating have been very fragmented. The HBP is driving integration of these different contributions.

During the Ramp-Up Phase, the HBP will collect strategic data, develop theoretical frameworks, and perform technical work necessary for the development of six Information and Communication Technology (ICT) Platforms during the Operational Phase. The ICT Platforms, offering services to neuroscientists, clinical researchers and technology developers, comprise Neuroinformatics (a data repository, including brain atlases and analysing tools); Brain Simulation (building ICT models and multi-scale simulations of brains and brain components); Medical Informatics (bringing together information on brain diseases); Neuromorphic Computing (ICT that mimics the functioning of the brain); and Neurorobotics (allowing testing of brain models and simulations in virtual environments). A High Performance Computing Platform will support these Platforms.

1.2 HBP Subproject 11: Applications

The overall goal of Subproject 11 is to prepare, evaluate and test the early applications of the six HBP Platforms. The Subproject is structured into three Work Packages, covering applications of the HBP Platforms in Future Neuroscience, Future Medicine and Future Computing.

The six HBP Platforms aim to provide unique capabilities for carrying out research that would not be otherwise possible in the fields of neuroscience, medicine and computing. The Platforms are scheduled to become operational at the end of the Project's Ramp-Up Phase (Month 30), so the main Applications work can only take place afterwards, during the Operational Phase, which will be governed by the HBP Framework Partnership Agreement (FPA). This will run from Month 31 to Month 120. When the HBP was conceived, it was clear that a set of early applications ideas should be developed in parallel with the building of the Platforms and that these early applications ideas should make use of preliminary versions of the Platforms' software and hardware.

1.3 Purpose of this Document

This report will describe progress in designing the experiments and creating the set-ups required for the HBP's planned Platform applications in neuroscience, medicine and computing.

1.4 Structure of this Document

The remainder of this chapter provides an SP-level overview, highlighting the SP's main accomplishments and problems encountered in the period M1-M12. Subsequent chapters look at accomplishments and in issues within individual components of the SP.

The Annexes present in tabular form what the Subproject planned to achieve in this period and what it actually achieved, including the Subproject's Scientific Key Performance Indicators (SKPIs).

- Future Neuroscience (WP11.1)
- Future Medicine (WP11.2)
- Future Computing (WP11.3)
- Scientific Coordination (WP11.4)

1.5 Overview of Subproject 11 Achievements

The Subproject delivered its comprehensive research plan (its first Deliverable) rather later than the planned deadline of Month 6. This delay was caused by the late, but necessary, integration of the Competitive Call groups that joined SP11 only in Month 7. Of the seven Milestones scheduled for achievement in the period Month 1 to Month 12, only 2 were not achieved within that timeframe (see next chapter).

Collaboration with other Subprojects is a key achievement for SP9, as it has to design and implement application cases for most of the HBP Platforms (each of which is built by a separate Subproject). SP11 performs particularly well in this activity. In Future Neuroscience, tight integration with SP10's Neurorobotics Platform and the cognitive neuroscience research in SP3 led to the design of an application experiment testing the Weber-Fechner Law and a set of integrated brain-body control benchmarks. In Future Medicine, achievement of the "Standardised description format for biological signatures of brain diseases" Milestone is a very big step towards the use of the Medical Informatics Platform (SP8). Again, cross-Platform cooperation has been the key to progress. In Future Computing, the integration in the latter part of Year 1 of the five new Applications Tasks which worth with the Neuromorphic Computing Platform (SP9) was a very important achievement. The new Tasks participate in all SP9 meetings and helped create the SP9 Specification Deliverable (D9.7.1). As such, they are now well informed about the performance and tools offered by the Neuromorphic Computing Platform.

Monitoring and quality control worked very well within the three research areas, but it has been a challenge to coordinate this across Work Packages 11.1, 11.2 and 11.3. Their work areas are still rather unrelated but it is expected that this will improve as the Platforms get closer to commissioning.

1.6 Overview of Subproject 11 Problems

Subproject 11 has not yet been able to reach Milestones 208 and 315. In the case of Milestone 208, this was due to late receipt of funding, which in turn delayed recruitment of staff. We now expect to reach this Milestone in May 2015. For Milestone 315, the number of network architectures that had to be evaluated was larger than originally planned. It is expected that a final decision on the network architecture to be implemented on the neuromorphic hardware will be made during the first few weeks of Project Year 2.

Neither of these delays risks jeopardising Subproject 11's goals or overall plan for the Ramp-Up Phase.

1.7 The Next Six Months for Subproject 11

During the first half of Project Year 2, the teams working on the Applications Use Cases will all complete their planning and prepare to implement them on the appropriate HBP Platforms: the Neurorobotics Platform for Future Neuroscience, the Medical Informatics Platform for Future Medicine and the Neuromorphic Computing Platform for Future Computing. These "host Platforms" will play key roles in the implementation of SP11 Applications in the first six months of Project Year 2.

2. Future Neuroscience (WP11.1)

2.1 Future Neuroscience: Overall Goals

The objective of this Work Package (WP) is to provide an initial demonstration of the value of the Neurorobotics Platform (NRP) for experimental cognitive neuroscience and to provide operational feedback on the design of the Platform. In the Ramp-Up Phase, WP11.1 will use the initial capabilities provided by the NRP to perform proof-of-concept simulation-based research into the multi-level brain mechanisms responsible for visual perception. In addition, this work may contribute to that of SP3's Task 3.1.2 (Understanding the circuits linking perceptions to actions, led by Martin Giese at Tübingen University in Germany). More broadly, this WP aims to demonstrate that the cognitive neuroscience community can use HBP technologies to enhance their scientific progress.

2.1.1 Psychophysics of Perception: the Weber-Fechner Law (T11.1.1)

This Task utilises the models, frameworks, and ideas of the HBP to connect theories of brain computing to psychophysical data that measure human performance. In close cooperation with T11.1.2, it is also exploring a methodology for benchmarking computational neuroscience models, by explicitly comparing specific psychophysical data against the visual pathway. In a sense, the goal is to provide input to the HBP model systems and measure outputs from those model systems for comparison with known properties of biological counterparts. The work is guided by two overarching ideas:

1) The neurophysiological properties of the cortical column identified by the HBP provide some guidance about which kinds of computations can be performed in cortical circuits for visual perception.
2) Given the uncertainties about the properties of the cortical circuits (e.g. the model is derived from data gathered from the mouse barrel cortex), it is appropriate to investigate behavioural data closely.

2.1.2 Integrated Brain-Body Control Benchmarks (T11.1.2)

This Task will study brain capabilities such as perception, attention, coordinated movements, core knowledge, spatial cognition, motivation, emotions and consciousness, It will also look at how these capabilities can be replicated in behavioural experiments using a robot. Quantitative data on behaviour and cognitive performance will serve as benchmarks for the validation of brain simulations and simplified brain models. Specific behaviours will be studied in closed loop systems linking brain simulations to simplified robot and / or virtual environments. T11.1.2 will provide benchmarks for the long-term goal of replicating classical cognitive neuroscience experiments using neurorobots and will perform comparisons between artificial and living animals. The results of these experiments will help validate brain/cognition models, and explore the structure and function of cognitive abilities for use in robots and other devices. More specifically, this Task will:

- Collaborate with the Cognitive Neuroscience Subproject in designing and performing cognitive neuroscience experiments.
- Initially use realistic, physics-based simulations and virtual robots to perform the experiments. After the Ramp-Up Phase, it will transition to physical robots.
- Support the building of experimental setups of increasing complexity, starting with simple experiments that explore and validate cognitive models of perception.

2.2 Future Neuroscience: Main Achievements (Revised Information)

The objective of Work Package 11.1 is to provide an initial demonstration of the value of the Neurorobotics Platform for experimental cognitive neuroscience, and to provide feedback on the design of the Platform. In the Ramp-Up Phase, WP11.1 has so far used the initial capabilities provided by the Neurorobotics Platform to perform proof-of-concept simulations. Only in the second phase of the HBP will the scope of these investigations be broadened, to match experimental investigations by the Subproject into brain function and cognitive architectures. Progress in W11.1 has been monitored by WP11.4.

The TUM team (*Technische Universität München – P53)* worked in close collaboration with SP10 to make sure that the development of the first close-loop experiment was well aligned with the development of the Neurorobotics Platform. At the beginning, a student jointly supervised by TUM and UGR (*Universidad de Granada – P58)* worked on retinal models for Task 11.1.2. TUM also worked on a physical, musculoskeletal platform based on variable stiffness joints with sensors, links, and motors. In collaboration with UGR, a real-time model of the cerebellum was selected as a motor controller. It was then ported to PyNN, so that it could be used as front-end for neuronal simulators including the SpiNNaker board, which allows for future connections to several other Tasks within this SP, and also connects to SP9.

All groups worked on the development of Milestone MS201, which developed an initial experimental design for testing the models for the Weber-Fechner law of visual perception. The planned design places the models within a simulated environment, with closed-loop interactions between the environment and the neural models. A document describing the experimental design has been uploaded to EMDesk. We encountered some difficulty producing benchmarks for the as-yet-unfinished Neurorobotics Platform, but we anticipate catching up when the Platform is released. In the meantime, Task 11.1.2 members worked on motor models that will promote the closed-loop experiment being integrated into the Neurorobotics Platform.

The Neurorobotics tool chain was successfully tested via the development of a "Braitenberg vehicle" experiment. This test demonstrated the interplay of different virtual robots (a Husky robot and a Lauron robot) with a virtual environment that includes objects such as bookshelves and pictures on the wall. A neural network model was calibrated to realise the brain-body transfer function. This work was done in close collaboration with SP10, especially WP10.5.

The next major task is to connect the retinal and cortical models into a single system, and then to integrate that combined model into the simulation environment being created by SP10. The resulting system will then be ported to various computing platforms being developed by the HBP (e.g., supercomputers, SpiNNaker).

2.3 Future Neuroscience: Main Problems

Areas where WP is behind

Although development of the retinal and cortical model has proceeded well, we have faced technical challenges connecting the two models.

Why behind

It has proven difficult to install the necessary libraries in a variety of systems. We have tried three different systems, but encountered problems each time (not always with the same library).

What is being done

The most recent unsuccessful efforts suggest that computer memory may be an issue (but this is not yet established). Additional memory has been ordered.

Likely impact

Because of this technical problem, we are a bit behind on those KPIs that require the integrated models. We believe the problem can be solved, and the impact is likely to be only be a few months delay in those KPIs.

2.4 Future Neuroscience: The Next Six Months

Milestone achievement

We anticipate that we will achieve MS202 "Robot, environment and experiment implemented" by completing a first implementation of all components (virtual robot, environment and experiment).

Other major accomplishments

Refinements of retinal and cortical models (SKPIs 11.1.1.4-11.1.1.7, 11.1.1.10 and 11.1.1.11). Exploration of models to account for psychophysical data (SKPIs 11.1.13-11.1.1.20). Exploration of models and cortical column statistics (SKPI 11.1.1.12). Integration of the retinal model with the simulation environment being developed in SP10 (SKPI 11.1.2.5). Development of a suitable set of benchmarks as contribution to MS202 (SKPI 11.1.2.3).

Work that can start now

All of the above.

Changes

None.

Major problems

Meeting Milestone MS202 requires integration with the Neurorobotics Platform. Our estimates of accomplishments over the next six months assume that the Platform will become operational according to schedule. Should that not be the case, we will proceed with simpler simulations that do not include a closed-loop environment.

3. Future Medicine (WP11.2)

3.1 Future Medicine: Overall Goals

SP8's Medical Informatics Platform will federate imaging, genetic and other clinical data currently stored in hospital/research archives and databases, but unavailable to the wider medical research community. SP8 will also provide tools for epidemiological exploration, numerical and statistical analysis, data visualisation and data mining. WP11.2 will use these resources to identify unique biological signatures of brain diseases.

With advanced computing power and data mining, nearly exhaustive searches of a data space can be performed to identify sets of rules that describe homogeneous populations, to explain their biological data and to predict the pattern of symptoms. Biological Signatures of Diseases are the results of a continuous dynamic data mining process of clinical data in local data sources. These will be used for diagnosis, more accurate prognosis and new approaches to drug discovery for the development of new medicines. The biological signature of brain diseases will form the basis for a new multi-dimensional brain disease space, facilitating scientific investigation and permitting personalised medicine. The advantage of disease signatures is that they are based on mechanistic, deterministic and predictive rules, as opposed to purely descriptive (phenomenological or clinical) features.

3.1.1 Biological Signatures of Diseases (T11.2.1)

This Task will rely on data accessible through the Medical Informatics Platform, including data on the longitudinal study of the large cohort of 'control' Alzheimer's patients. Even with common Alzheimer's disease (AD), the clinical syndromic diagnoses are wrong in 20% cases. The current study is designed to show that, with modern mathematics, powerful information technology and a large data set, it is possible to identify homogeneous groups of patients, characterised by a set of parameterised latent causes, which constitute what we choose to call "disease signatures". There is a great amount of uncertainty regarding the accuracy of diagnostic classification in the early stages of AD, because of the underlying heterogeneity in etiologies leading to similar phenotype. To explain the observed heterogeneity we use rule based clustering algorithm and identify homogeneous subgroups of patients. The hypothesis is that such subgroups are due to the same underlying causes.

3.2 Future Medicine: Main Achievements

WP Milestones achieved

The Milestone MS204 "Standardised description format for biological signatures of brain disease" has been achieved. The on-going study of Alzheimer's disease phenotype using machine learning methods and brain pathology has identified several subgroups or subtypes of AD and therefore defined the first standard description of the biological signatures of brain diseases.

Other WP progress

The Function (11.2.1.1) providing a "Description format for the biological signatures of the disease" has been achieved. The following KPIs have been achieved: 1) Identify multimodal clinical data, 2) Data pre-processing, 3) Data aligned, and 4) Feature selection. Our work was based on the observation that previous neuroimaging-based classification can predict Alzheimer's disease pathological diagnosis. We generated a classifier using Support Vector

Machine (SVM) and applied it to predict pathological diagnosis for the living subjects from the Alzheimer's Disease Neuroimaging Initiative (ADNI) cohort that have been clinically diagnosed. Among the clinical healthy control (HC) group, we identified a group of subjects at risk, because they are showing disease-specific atrophy. We confirmed that their progress from this state to abnormal cognition condition is significantly faster in the follow up study than the remaining control subjects. Previously, we have characterised the subgroups in term of brain anatomy. We extended this approach to cognitive measures and genetic features. The results of this study were presented at the Neurological International Conference (Istanbul-Turkey, May 2014). A manuscript for a peer-reviewed article is in preparation.

The work is progressing on the next Function (11.2.1.2) "Informatics based model for generating biological signature of a disease". In this task, we aim to test and benchmark other type of algorithms for building the model of disease. An important aspect is to design a method that can select the best features. For this Function a neural networks algorithm based on stacked auto-encoder were used to classify grey matter images of AD and HC subjects. Neural networks with multiple hidden layers can be used to classify input vectors. In this work, we proposed two different ways of feature selection for training stacked auto-encoders for the classification of AD versus HC subjects. As opposed to previous studies on this topic – most of them aiming to optimise the classification accuracy – no initial analysis of the raw scans was performed to preselect significant regions of the brain. On the contrary, features were selected in such a way that full scans do not have to be compared, which open perspectives when privacy issues might occur. Three hidden layers containing 800 + 500 + 50 neurons were used. This configuration was chosen out of 27 possible combinations as it gave the best classification result for a random subset of 64 subjects. The number of epochs was set to 200 and the batch size to 100. We used a 15-fold setup in order to test all subjects, resulting in a training set of 600 subjects or 60,000 images. Training the network takes around 6 hours on a normal computer.

The results show that low level features, extracted from the scans without any preliminary knowledge, combined with a fairly easy deep learning scheme has promising classification potential. Using volumes of grey matter ROI defined by a brain atlas, HC subjects are correctly classified in 75% of the cases and only 7% are consequently misclassified. The performance accuracy is lower for AD classification (57% correct vs. 19% misclassified), but the misclassified subjects actually show structural properties of the HC subjects (almost no atrophy in temporal and hippocampal areas). Additionally, the Mini Mental State Examination (MMSE) scores of misclassified AD subjects are significantly higher than the true AD subjects, which means that this could be a subgroup needing a different label referring to a mild or early stage of the disease. Apart from having the capability to classify single subjects (which is not possible using state-of-the art statistical methods such as Voxel-Based-Morphometry), additional conclusions such as the definition of subgroups or finding brain regions affected by a disease.

Collaboration

WP11.2 works closely with SP8 to increase the volume of data for the algorithms developed and tested. Data are obtained from large cohort research studies.

Results from combination of data and ICT

T11.2.1 collaborates with SP8 in the design of the Medical Informatics Platform. The data mining algorithms for the identification of biological signatures of diseases have been specified in the Deliverable D8.6.1 of WP8.6 (technical specifications and implementation within the MIP infrastructure). The document describes the layers that will support the applications provided by the Platform to researchers. From basic statistical description to

advance analytic and machine learning tools the Platform will help research to query data and explore the biological signature of brain diseases.

Internal monitoring and quality control activities

Work is discussed in a regular joint meeting with SP8.

Outstanding contributions to WP work

An important achievement is that we show that computer aided diagnosis of AD has proven to be a promising method of early detection, an important condition for a more effective treatment of the disease.

3.3 Future Medicine: Main Problems

None.

3.4 Future Medicine: The Next Six Months

The next Milestone MS205 "First draft informatics-based model generating a biological signature of a disease" should be achieved as expected.

In the next six months, work will be carried out to achieve a more accurate version of "Informatics based model for generating biological signature of a disease". The work will consist of the implementation and testing of different algorithms, configuration of the model and selection of algorithms.

T11.2.1 is now testing algorithms for data mining. Thanks to the automation of the pre-processing steps, new data can be easily added to the system. However, more computing power will be required to extract the relevant features for the data-mining algorithm, in particular those based on "deep learning" and neural networks.

Any changes you will need to make to the WP's original Ramp-Up Phase plan

None.

Major problems that the WP might encounter & your plans to address them

None.

4. Future Computing (WP11.3)

4.1 Future Computing: Overall Goals

Applications in Future Computing will prepare for the implementation of neural circuit models on the hardware systems provided by the Neuromorphic Computing Platform. The initial network architectures are compatible with the constraints given by the HBP hardware architectures, as described in the SP9 Platform Specification Document, and will all be formulated in the network description language PyNN. The HBP will offer a unique Neuromorphic Computing Platform to explore the computational capabilities of spiking neural networks. Two complementary hardware systems are being constructed, a "many-core" model in Manchester (NM-MC-1) and a "physical model" in Heidelberg (NM-PM-1).

The NM-MC-1 system will consist of 0.5 million ARM cores, with 18 cores per chip. The cores provide integer operation capability and each chip has six bi-directional links with a bandwidth for 6 million spikes per second per link. Networks simulated on this system will operate in real time. The NM-PM-1 system will be a physical model based on analogue neurons and synapses, with binary, asynchronous and continuous time spike communication. NM-PM-1 will provide 20 individual wafer systems, replicating a total of 4 million neurons and 1 billion synapses. Networks simulated on this system will operate 10,000 times faster than real time. Both systems are expected to commence hardware operation towards the end of the HBP's Ramp-Up Phase.

4.1.1 Neuromorphic Data Mining Systems (T11.3.1)

Very large, open-ended data streams are generated in industrial applications by real-time transactions systems in various industries like the retail sector, by utility companies, in production facilities, by sensors in surveillance systems, to name but a few. In contrast to traditional data sets, stream data flows in and out of a computer system and it may be impossible to store the entire stream due to its large volume. In cases where the data are actually stored, multiple scans become very expensive and single-scan analysis methods need to be applied [1]. Within this continuous stream of data, the identification of spatial-temporal patterns is probably the most important task and despite the fact that this has been an active research field for many years no general-purpose solution to the problem is currently available. Therefore, the goal of this Task is to contribute to this active research field by the adaption/modification of existing algorithms or by the development of a novel algorithm.

Requirements for the algorithm include the ability to detect hidden patterns and causal relationships. A software implementation should be developed as a proof-of-concept with non-spiking neurons. In a second step, the algorithm will be extended to a spiking network system. Finally, the implementation will be migrated to the Neuromorphic Computing Platform provided by the Human Brain Project (HBP). The final proof point will be the ability of the algorithm to handle real-world business relevant data sets and to demonstrate the validity of the Neuromorphic Computing Platform to implement a data mining system.

The overarching goal of the work is to demonstrate the validity of the Neuromorphic Computing Platform as a tool for producing and prototyping cognitive devices and systems outside the realm of biology, in particular to validate the possibility to develop cognitive business information systems.

4.1.2 Port CABot3 to Neuromorphic Chips and Extend (T11.3.2)

This Task will contribute to the HBP by providing an extensible agent that can be used in a 3D environment, and by providing advances for the HBP's neuromorphic systems to learn tasks normally managed by Cornu Ammonis (CA) areas in the Hippocampus. The agent will be useful for researchers to use during and after the Ramp-Up Phase. It will be a working embodied agent in a simulated environment implemented entirely in simulated neurons, providing a modifiable early link in the Project between robotic systems, cognitive architectures, brain data and neuromorphic hardware. It will provide existing modules for neural language, vision, planning and action, which have a reasonable degree of modularity. Improved models of CA learning will provide insights into the theoretical problem of concept formation in neural systems. This work will be linked to the environment and psychology, as well as known and posited neural behaviour, leading to a significant impact in both the short and longer term. The objectives and approach for T11.3.2 are summarised below:

Neuromorphic Computing Systems Objectives:

- Implement novel computing paradigms using an embodied agent in spiking neurons.
- Develop generic circuit concepts of spiking neurons through improved cell assembly models.

Operational Objectives:

- Advance the Neuromorphic Platform by implementing simplified versions of brain models.
- Advance theoretical foundations by furthering generic CA models.
- Advance the Brain Simulation Platform by building point models for simulating brain areas.
- Advance cognitive architectures by building neuro-cognitive models to extract principles.

4.1.3 Exploiting of Feedback in Ultra-Fast Spiking Visual Architectures (T11.3.3)

This Task will port architectures for spike-based visual object recognition to PyNN and preliminary simulations on the HBP Platforms will be started. The feed forward architectures will be similar to those simulations recently reported by IMSE [2] where recognition was performed with delays in the range of 1-2ms. Exhaustive optimisations will be performed, both on system parameters (such as synaptic values, neurons thresholds, delays, etc.) and structural properties (such as number of layers, number of feature extractors per layer, number of neurons per feature extractor, etc.). Gradual Attentional Feedback will be added to the architectures to carefully assess performance variations in object recognition while degrading the visual stimuli. Additional hardware interfacing between Address-Event Representation (AER) sensors and processing modules with the HBP systems will be conducted. Finally, mismatch impact on performance will be analysed.

4.1.4 Spiking Associative Networks for Neuromorphic Computing Systems (T11.3.4)

The overall aim of this Task is to explore the computational capabilities of spiking neural networks for associative memory (SAM: spiking associative memory) with the help of the two complementary HBP hardware systems in Heidelberg and Manchester. Neural associative memories can be seen as models for local networks in the cerebral cortex [3] and are closely related to assembly theory [4]. Furthermore, they offer efficient methods for fault tolerant information retrieval in technical cognitive systems with possibly better

performance than classical algorithms of computer science-like content addressable memories, search-trees, or hash-tables[5]. We start with the classical binary associative net with binary neurons and synapses proposed by Steinbuch[6], and theoretically analysed by Willshaw *et al*. [7] and Palm[8] which is in the following referred to as the BiNAM (binary neural associative memory).

The Task will have two parts (modelling SAMs with PyNN and HBP hardware; SAM design space exploration). First, we will focus on the implementation of spiking associative memories (SAMs) in the simulator-independent network description language PyNN and their mapping onto the HBP hardware. The first SAM version will be based on the leaky integrated-and-fire neuron as implemented in the Heidelberg hardware[9] and the spiking associative memory based on BiNAM as implemented by Knoblauch *et al*[10] [11]. Step-by-step, we will extend this model, e.g. with non-binary synapses, inhibitory neurons and more complex feedback structures, and will study the effects on SAM behaviour quantified by the performance measures (storage capacity, retrieval time, energy consumption, and robustness). For functional verification, we will use a software simulator (e.g. Neuron, Nest, PCSIM). In order to be compatible with the HBP hardware suppliers, we will coordinate the selection of the simulator with Heidelberg and Manchester. The next step will be the mapping of the SAM to the virtual hardware simulator of the Heidelberg group, which is a detailed simulation of the final wafer-scale hardware system implemented in C++/SystemC [12]. The virtual hardware offers an early modeller's perspective onto the capabilities of the future wafer-scale system. The last step will be the mapping to the hardware systems (Heidelberg, Manchester) as soon as they are available for this Task. For a smooth transition from simulation to real hardware we intend to use the HICANN prototype.

4.1.5 Asynchronous Computational Retina (T11.3.5)

The Task aims to develop a pure event-driven visual computation approach that will use precise timing mechanisms to design new computation techniques in visual processing. The Task will produce a full event-driven visual processing system linking a neuromorphic retina directly to the SpiNNaker by an Asynchronous Event Representation (AER) bus. The architecture will allow for the first time the development and implementation in real-time of new visual event-driven computation techniques. Event-driven early vision models and 3D stereovision that will use precise timing mechanism, will be implemented into SpiNNaker board and fed directly by the outputs of a neuromorphic retina, ATIS (Asynchronous Time-based Image Sensor), developed by the UPMC's research team. The consistency and the robustness of the implemented models and algorithms will be constantly analysed through extensive evaluations carried out all throughout the Task.

In this Task, we propose to feed the event-driven retina sensor's outputs directly into SpiNNaker, in order to make use of the temporal dynamics of these sensors. This will allow us to build a neuromorphic sensory architecture able to process and interpret visual information at the moment they are acquired thus keeping the relative time between events. To reach this goal, we are using the ATIS silicon retina[13].

The proposed approach will be organised as follows:

Interfacing the ATIS with SpiNNaker: This primary task is mandatory to allow SpiNNaker to process the retina's inputs as directly as possible from the silicon retina. Both devices are already using the same event-driven data format (AER). However, their designs do not include natively their connexion. The interfacing will also be extended to allow the integration of two ATIS in order to get the basis of an event-driven stereovision system.

The event-driven computation paradigm for artificial and asynchronous retina model: This Task's objective is to implement the retina model developed in[14], in which the behaviour

of ganglion cells is reproduced on the SpiNNaker board. The higher temporal and spatial accuracy of the ATIS allows outlining the impact of precise timing: the event-driven optical flow developed in [12] and the event-driven 3D reconstruction algorithm that triangulate points from the stereovision system[15] [16]. The implementation of these algorithms onto SpiNNaker serves the purpose of the asynchronous processing of the flows of events at the rate of their arrival using an event-driven methodology.

4.1.6 Implementing a Spiking Classifier Network on HiCANN (T11.3.6)

The overarching goal for this Task is to implement spiking neuromorphic networks for multivariate data classification, which exploit the properties of large-scale neuromorphic systems and solve real-world computing problems. Our networks harness the massively parallel architecture of brain circuits to efficiently process high-dimensional data. Our network design is inspired by the olfactory system, mimicking the general blueprint of parallel feature encoding, processing and classification in feed-forward networks with lateral inhibition. Achieving our overall goal entails addressing the following objectives:

- To develop a scalable implementation targeting the large-scale system provided by the Heidelberg group.
- To implement this solution on the Heidelberg system and test it on large multivariate classification problems.

4.2 Future Computing: Main Achievements (Additional and Revised Information)

4.2.1 Neuromorphic Data Mining Systems (T11.3.1) (Additional Information)

The evaluation of potential application scenarios and the corresponding algorithms (11.3.1.1) proceeded as described in the Research Plan D11.4.1. After an expanded analysis of Use Cases of different areas, the current plan is to focus on the three scenarios coming from complementary application domains. The first case is related to optimisation of data base management systems, the second deals with forecasting of business data and the third case will most likely come from the IoT (Internet-of-Things) domain. During the evaluation of suitable network architectures many different algorithms have been reviewed. The analysis was performed with respect to three main criteria: (i) possibility of an implementation on the neuromorphic hardware (ii) complexity of the algorithm to ensure a staged implementation and testing approach, start simple and increase the complexity stepwise (iii) suitability for a business relevant use case. A short list of network structures and algorithms has been compiled. The most promising candidate seems to be a model of the Cerebella published decades ago and the corresponding algorithm. It is capable to approximate any non-linear function, exhibits fast learning, reasonable generalisation capabilities and robust noise tolerance. The final decision on which network structure to implement will be made in the next couple of months.

Application domains considered as potential use cases range from demand forecasting in retail stores, fraud detection in insurance claims, supply chain optimisation problems and context-aware user guidance, to database management systems. In all these areas, SAP has successful products on the market, and deep technical and business expertise.

While SAP has extensive machine learning knowledge, experience in the field of spiking neural networks currently resides mainly in academia and research institutes. IBM has developed its own Neuromorphic chip, which is non-spiking. Based on our background, we started looking at ‘classical’ neural network structures, keeping in mind the applicability to our use cases and the possibility to create spiking versions running on the hardware. Initially, we focused on various versions of associative memories. In particular, recent

development related to the HTM (Hierarchical Temporal Memory) framework developed by Jeff HAWKINS and Dileep GEORGE was reviewed. Due to the complexity of the models (HTM requires several hierarchies and is still non-spiking), and considering the time and resource constraints, the decision was made to concentrate on simpler models of associative memories. Recently, IBM and Numenta set up a joint project with 100 people to test the HTM algorithm, and to look at its use in a potential hardware implementation (http://www.digitaltrends.com/computing/ibm-creates-a-research-group-to-test-numenta-a-brain-like-ai-software/).

The Cerebellum model and the corresponding CMAC algorithm published by Albus in 1975 represent a less complex version of an associative memory function. This algorithm combines several attractive features and is still simple enough to be implemented on the available neuromorphic HW platforms. Hardware implementations of the CMAC algorithm exist, but to our knowledge most of them are based on a look-up version of the algorithm, implemented in FPGAs, and are not based on the neural structure of the model. We believe that a neural implementation of the CMAC algorithm based on spiking neurons is feasible, and that the implementation would exhibit interesting features to test our use cases. This has been demonstrated in software simulations, and is described below in more detail for one of the use cases.

The broad range of reviewed algorithms led to Milestone MS209 (network architecture evaluation) being delayed. A network structure was decided in the meantime, and the Milestone was reached. This delay will not have any detrimental effect on the overall progress of the Task.

4.2.2 Port CABot3 to Neuromorphic Chips and Extend (T11.3.2) (Revised Information)

The agent was successfully ported to SpiNNaker in March 2015. The neurocomputational model had to be rewritten from scratch. We chose to use a standard neural model that ran on SpiNNaker and Nest. Initially, we tried an Izhekivch model, but found problems with the neurons spiking when inhibited. We then switched to a Brette and Gerstner model, and proceeded. Initial versions of the subsystems were developed to run in Nest, and then using the SpiNNaker emulator. Eventually, a 4-chip SpiNNaker board was used for an agent. A virtual environment was developed, and interaction between the board and the environment was managed; at this stage, Nest was no longer supported for the agent, though we hope to reintegrate it in the near future. Once we got a 48-chip SpiNNaker board, we successfully completed the port. The agent takes natural language commands to set goals. The main input to the agent is the visual scene as viewed from the avatar's camera. The agent is parameterised and can take the visual scene in as an NxN grid; the agent works with 30x30 input for the 48-chip board, and is less successful with 25x25 input on the 4-chip board. Larger input can cause IO problems. The agent also receives user commands in the form of words, and outputs actions. The goals interact with the current environment via the planning system. Simple commands such as "Turn right." or "Move forward." are successful. Compound commands, e.g. "Move left." (a combination of turn left and move forward) are successful. Environmentally sensitive commands (e.g. "Turn toward the pyramid." or "Move to the pyramid.") are usually successful, though they depend on the object being in the visual field.

The most complex commands involve the cognitive map. The agent makes a simple cognitive map of the environment. The environment consists of four rooms connected by corridors. Each room contains a unique shape, a red or blue pyramid or stalactite. The command "Explore." causes the agent to move around the environment and form a room-to-shape map. This is tested by a command such as "Go to the room before the blue pyramid." This causes the agent to move clockwise around the environment, and stop in

the appropriate room. This shows that the map has been learned. This full process, explore and go, is successful roughly 90% of the time when the agent starts with the first item in the visual field. In the above example, successful means that it works every time we try it. We are hoping to do a full evaluation this summer and write a corresponding paper.

Since the completion of the initial port, we have made progress on several internal waypoints. We have made some progress with Cell Assemblies (CAs) that persist for reasonable lengths of time. We have also made some progress with a cognitive model of categorisation. We have developed new synaptic and neural models in Nest. We have started with the ESS (HICANN) emulator, but have yet to get it working. The internal waypoint of having the model run on HICANN has not been met because of this, and the lack of the actual chip.

4.2.3 Exploiting of Feedback in Ultra-Fast Spiking Visual Architectures (T11.3.3)

In this period, we achieved the Milestone MS310 "Initial sensor interfaces operative", corresponding to Function 11.3.3.4 (within Step B). We successfully implemented an AER interface between our AER in-house hardware and the SpiNNaker hardware, capable of communicating through commercial serial SATA connectors using commercial Spartan6 FPGAs. We have tested transmitted event rates of up to 62.5Meps (Mega events per second) with 32bit events per LVDS link. This rate is for board-to-board connected with SATA. If one of the boards connects to a SpiNNaker interface the measured transmission rate drops to 1.1Meps. This is mainly due to the nature of the SpiNNaker asynchronous protocol, which is meant for very low power and is composed of a complex sequence of cycles, thus reducing the throughput.

There has also been progress in parallel on providing an initial PyNN description of a Convolutional Neural Network for object recognition, and testing it on SpiNNaker. This corresponds to Function 11.3.3.1 (Step A1) "Description of multi-layer vision system in PyNN". We faced some difficulties because the supporting software (pacman) has been changing and had to readapt items. We decided to send a student over to UMAN (but paid by another project) to advance on this aspect. Progress is satisfactory at this moment.

Collaboration with UMAN have been important for progress, as stated above.

We have now hired one person, paid by HBP, who started on 1 September 2014.

4.2.4 Spiking Associative Networks for Neuromorphic Computing Systems (T11.3.4)

The SANNCS project started in April 2014 and was presented at the HBP SP9 meeting at Sabanci University, Istanbul (10 April 2014). During this meeting, important coordination and cooperation discussions took place in order to integrate the SANNCS smoothly into the Neuromorphic Computing Division of HBP.

In the first six months of the Project, we successfully implemented a generic software prototype of a spiking associative memory (SAM) in PyNN based on leaky integrate-and-fire (IaF) neurons. The SAM has a single layer structure and the synapses of the IaF neurons are binary. This prototype is now used to explore the effects of parameter variations on the behaviour of SAMs. For a more detailed analysis of the parameter influences, we started to implement the specified programme package for automatic generation of SAM benchmark data sets and the software tool for comprehensive reporting of the performance evaluation based on the simulation outputs (spike trains). First results of our research were discussed with our Partners at the SP9 meeting in Berlin (24 July 2014) and the HBP Summit in Heidelberg (28 September – 1 October 2014).

We contributed to the definition of the D11.4.1 Deliverable and specified scientific KPIs for our Task. Furthermore, UNIBI participated in the HBP Framework Partnership Agreement (FPA). We also joined the special interest group on neuromorphic hardware benchmarking (SP9), led by Anders Lansner, KTH Stockholm. Together with the groups from KTH Stockholm and UHEI Heidelberg, we contribute to the definition of benchmark sets for spiking associative memories.

4.2.5 Asynchronous Computational Retina (T11.3.5)

The current period includes Milestone MS315, which was postponed from M12 to M19 (see sections 4.3.5 and 4.4.5). Work has been done towards the achievement of this Milestone, corresponding to the completion of Functions 11.3.5.1, 11.3.5.2, 11.3.5.3 and 11.3.5.4. Functions 11.3.5.1 and 11.3.5.2 relating to interfacing the ATIS camera to the SpiNNaker system is currently under final testing and should be completed soon. The stimulation platform of Function 11.3.5.3 has been assembled and is in the process of software design to start testing its performances. Function 11.3.5.4, the database platform, has been physically set up and is waiting for Function 11.3.5.3 to move further.

In addition to the SpiNN-3 SpiNNaker board (embedding 4 chips) available for our work, we have just received from the SpiNNaker group in Manchester a SpiNN-5 SpiNNaker board (embedding 48 chips), which will increase the processing power available for implementation of Functions 11.3.5.5, 11.3.5.6 and 11.3.5.7. The interface developed for Function 11.3.5.1 has been tested on this new hardware and is fully functional.

Xavier Lagorce has also been recruited, working full time on the Task from September 2014 to August 2015.

4.2.6 Implementing a Spiking Classifier Network on HiCANN (T11.3.6)

The first Milestone for this Project is still pending. We recruited Alan Diamond to work full time on the Task from October 2014 until January 2016. So far, we have ported the spiking classifier network to a GPU Platform and analysed how the network scales to larger neuron numbers. We scaled the network to use ten times the number of neurons that the version that ran on Spikey used. This scaling allowed us to tackle novel classification problems, in particular an e-nose odour identification task.

4.3 Future Computing: Main Problems (Additional and Revised Information)

4.3.1 Neuromorphic Data Mining Systems (T11.3.1) (Additional Information)

The biggest challenge in the last 12 months was to identify applications, which benefit most from this new computing paradigm. The search for a “killer application” will remain an ongoing activity or the foreseeable future. The main problem is to understand the capabilities of the neuromorphic hardware and to identify/develop robust algorithms to address specific application scenarios. This will continue to be a challenge as the hardware evolves and much bigger systems will become available. The complexity of the algorithms and the corresponding network structures will increase, making it possible to address completely new application domains.

Our search of the application space is now determined by the following search parameters:

- Algorithm mappable to both types of neuromorphic hardware

- Clear differentiating value of neuromorphic hardware in a deployment scenario through either speed and/or power (precision has to be comparable to alternative software solutions).

For many of the originally envisaged applications, the differential value of the neuromorphic hardware compared to the complication of embedding the new hardware into the scenario has so far proved unconvincing. The remaining two scenarios we have chosen follow two extremes of the second differentiating factor — one goes for power efficiency in an embedded, distributed scenario with state-of-the-art precision, making it a good fit for the ARM hybrid, while the other goes for a demonstration of speed, playing toward the strengths of the analogue neuromorphic hardware.

4.3.2 Port CABot3 to Neuromorphic Chips and Extend (T11.3.2) (Revised Information)

We faced several problems in the initial port of the agent. These included a misstep using Izhekevich neurons, inadequate STP learning rules, problems with IO to the board and PyNN, Nest and SpiNNaker software changes. Our first problem was selecting a neural model. We wanted to use a standard model that ran on both Nest and SpiNNaker, and after a brief exploration chose Izhekevich neurons. Unfortunately, we found them difficult to work with because they spike rapidly when heavily inhibited, and our basic mechanism often uses inhibition. This caused an initial delay. Other delays were caused by the middleware changing, and incompatibilities between different engines. Finally, we spent a lot of time managing IO to the SpiNNaker board. Additionally, the initial CABot3 model (running in Java) uses short-term potentiation to bind in parsing; bindings are erased when the neurons do not fire. The Tsodyks-Markram STP synapse does not erase when the neurons do not fire. So, the parsing component does not use binding, but is merely a regular grammar parser.

4.3.3 Exploiting of Feedback in Ultra-Fast Spiking Visual Architectures (T11.3.3)

No Milestones have been missed at this moment, and for now we are not behind the expected work.

4.3.4 Spiking Associative Networks for Neuromorphic Computing Systems (T11.3.4)

The Project is on schedule. The first Milestone is planned for Month 19 (SAM implementation on virtual hardware).

4.3.5 Asynchronous Computational Retina (T11.3.5)

The Task has known some delays in the period covered by this report, with Milestone MS315 being postponed from M12 to M19 because Functions 11.3.5.2, 11.3.5.3 and 11.3.5.4 are currently behind schedule.

The hardware differences between the SpiNN-3 and SpiNN-5 SpiNNaker boards are forcing us to change our approach to Function 11.3.5.2. When the SpiNN-3 board has two SpiNN-links available to plug two ATIS cameras, the SpiNN-5 board has only one available (thus allowing direct connexion of only one camera) but also provides a new type of input connexions. We are working with the SpiNNaker team in Manchester (NM-PC-1) to adapt our designs to these new constraints (the SpiNN-5 board has been received in October 2014).

Due to delays in receiving the funding from the Project, we were unable to fill several internships, which had been planned for the summer period (April to August 2014). This

impacted Functions 11.3.5.3 and 11.3.5.4, which had less manpower than planned in this critical period. The late availability of these funds also resulted in delays in the availability of supplies and hardware platforms for these functions. Most of our suppliers for these sorts of items slow down their production in the summer period. As a result, these Functions were considerably delayed. Physical setups for Functions 11.3.5.3 and 11.3.5.4 were received in August 2014 and work on these Functions could only really start in September 2014, instead of April 2014.

4.3.6 Implementing a Spiking Classifier Network on HiCANN (T11.3.6)

This Task is on schedule according to the original Project plan. We anticipate that there may be problems with our upcoming Milestone "Implementation on HiCANN using adaptive neurons and a slightly higher neuron count" planned for MS318 (April 2015), due to its dependence on the availability of the HiCANN chip. However, contingency plans are already in place (see 4.4.6 below).

4.4 Future Computing: The Next Six Months (Additional and Revised Information)

4.4.1 Neuromorphic Data Mining Systems (T11.3.1) (Additional Information)

During the next six months, the implementation of the selected algorithm and network architecture in the hardware description language PyNN will be pursued (11.3.1.2). The goal is to have a running implementation in PyNN and to perform first tests on the neuromorphic hardware. In addition, an MSc student will join the team to support implementation and testing activities (11.3.1.3).

Work on demonstrating the superior execution speed and efficiency for the analogue neuromorphic hardware is progressing. After selecting a suitable algorithm that is known to be executable on the somewhat constrained environment of the analogue neuromorphic hardware, we proceeded to create a suitable training environment. The major obstacle here was that we needed a supervised training setup for our classifier and could not get a suitable training signal from the real target environment. This required us to recreate a software simulation of the target (not unlike the virtual worlds approach used for the robotics Work Package) that had to be detailed enough to exhibit the same complex nonlinear interactions as the real target hardware.

The early intermediate results show that we can learn reasonable approximations of the optimisation problems with the chosen algorithm. However, we cannot yet demonstrate sufficient efficiency gains compared to simpler heuristic approaches to warrant the far more elaborate hardware approach.

The next section describes SAP's efforts related to the use case "predicitve maintenance using (spiking) neural networks". This work started in the reporting period and is an on-going activity.

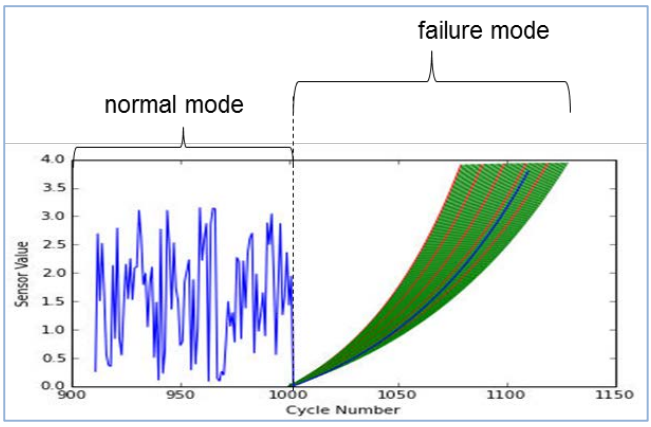

The problem posed by this use case is that of determining the remaining useful life (RUL) of a machine based on its sensor readings. Training data sets (sensor readings) for a limited number of breakdown scenarios are available

(supervised training), however generalisation to a wider range of previously un-observed sensor readings is required. The main motivation for this use case is to explore the possibility of moving application logic closer to the source of data creation (machine, embedded sensor) using the low power and high-speed capabilities/parallelism of the HW, and to significantly reduce the data traffic to the backend. During the breakdown of a machine, certain sensory readings show a typical behaviour, and numerous models are available to describe the damage propagation. The exponential evolution of the fault signal is common to all these degradation models. Assuming an upper threshold that describes an operation limit beyond which the machine cannot be used anymore, a health index H(t) can be defined.

Model data has been generated describing the failure mode and the normal mode. The CMAC network mentioned above was trained with these data sets for (a) a classification problem and (b) approximating the health function.

Preliminary results are shown below; the network could clearly distinguish between the two operation modes and in case of the failure mode, a reasonable approximation of the target function could be reproduced. It remains to be shown whether the described software simulation can be implemented on the HW, and whether the results show an advantage compared to conventional approaches.

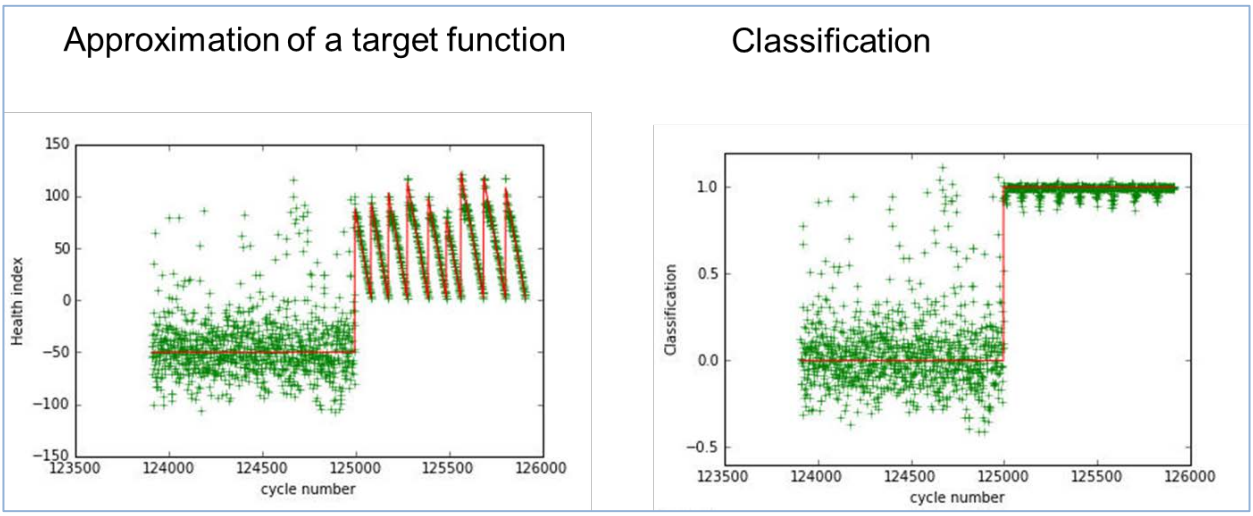

The work described summarises SAP's effort related to T11.3.1 for the first reporting period of approximately 12 PM, as planned in the DoW.

4.4.2 Port CABot3 to Neuromorphic Chips and Extend (T11.3.2)

We expect to have met all our internal waypoints in the next six months, and expect to have a working categorisation model. This will probably require new neural and synaptic models running on Nest and SpiNNaker. We expect to have HiCANN running, at least the emulator, and have the basic neural model running on it. We also expect to have the agent running on HiCANN, but this is less certain as there may be unanticipated differences. We expect to have plans cached away by the system, a simple form of learning. We expect to have CAs that persist for psychologically reasonable times.

4.4.3 Exploiting of Feedback in Ultra-Fast Spiking Visual Architectures (T11.3.3)

In the next six months, we will start focusing on the four main aspects of our scientific contribution:

(a) Perform parameter optimisations. This corresponds to Function 11.3.3.2 "Simulations with exhaustive optimisation", where we will adapt the PyNN description results that came out of Function 11.3.3.1 to vary critical parameters using some optimisation algorithm.

(b) Start analysing the effects of introducing feedback. This corresponds to Function 11.3.3.3 "PyNN of feedback architecture". The first job will be to adapt the PyNN description available from Function 11.3.3.1 to introduce feedback paths.

(c) Start introducing the effects of mismatch. This corresponds to Function 11.3.3.7 "PyNN for mismatch". Again, we will use the PyNN description available from Function 11.3.3.1 to add statistical random parameter variations to analyse mismatch effects.

(d) Start elaborating hybrid hardware, combining our AER hardware with SpiNNaker. This corresponds to Function 11.3.3.5 "Integrating multi-layer vision system into the neuromorphic Platforms". Here we will take advantage of the results from Function 11.3.3.4 on "initial interfaces characterisations", so that now AER sensors and AER processing hardware available in our group can interconnect to SpiNNaker using these initial interfaces. We will continue studying and characterising these interfaces, identify possible shortcomings, and try to improve them.

There are no specific Milestones in this period.

4.4.4 Spiking Associative Networks for Neuromorphic Computing Systems (T11.3.4)

During the next six months, we plan to implement another SAM in PyNN (Function 11.3.4.1). In parallel, we will improve our tools for automatic performance SAM evaluation (Function 11.3.4.2). In 2015, we plan to start with the automatic generation of benchmark data sets (Function 11.3.4.3) and mapping of the first SAM architecture on the virtual hardware simulator (Function 11.3.4.4).

4.4.5 Asynchronous Computational Retina (T11.3.5) (Revised Information)

Milestone MS315 was not achieved in Month 12, but in Month 19. Even so, progress was traceable along functions. Please find detailed information on this Milestone in Appendix A.

4.4.6 Implementing a Spiking Classifier Network on HiCANN (T11.3.6)

MS318 "Implementation on HiCANN using adaptive neurons and a slightly higher neuron count" is scheduled for April 201). Achievement of this Milestone is conditional on the availability of the HiCANN chip for PyNN users. As a fall-back, we are planning to port the network to the SpiNNaker Platform and Dr. Diamond will attend a SpiNNaker training event in January 2015. The SpiNNaker Platform will allow us to implement neuron counts of a similar order of magnitude as eventually planned for the wafer-scale system, as per MS319, due in M27 (December 2015).

5. Scientific Coordination and Support (WP11.4)

5.1 Scientific Coordination: Internal Meetings

Scientific coordination across this Subproject is quite a challenge due to the fact that the interaction between each SP11 Work Package and its associated Subproject (WP11-1 – SP10, WP11.2 – SP8, and WP11.3 – SP9) is significantly stronger than any connection between the SP11 Work Packages.

For this reason, coordination was reduced to informal updates on Deliverable organisation and collection of contributions, reporting duties, forwarding of review information, EC requests, etc. For this purpose, an internal mailing list was set-up and a few video-conferences were held.

Integrating the five new Partners into the Project was one of the main goals in the third quarter of the year. The new Partners presented their Task work plans at a physical meeting in Sabanci between SP9 and WP11.3, and a videoconference of all SP11 Task leaders of SP11.

This table lists meetings between SP staff.

Date	Description	Location	Participants	Comments
16.11.2013	Planning retinal model (SP11) Discussed properties of current simulations, plans for future development, and properties of the cortical model simulators. Agreed that NEF format for spike trains is good for everyone.	VC	Greg Francis (EPFL), Marc-Oliver Gewaltig (EPFL), Eduardo Ros (UGR)	
05.12.2013	Analysing retinal model output (SP11) Discussed the properties of the retinal data output provided by the UGR group. Considered what aspects of the output were most important for integration with a cortical model.	EPFL	Greg Francis (EPFL), Marc-Oliver Gewaltig (EPFL)	
19.12.2013	Analysing retinal model output (SP11) Discussed the properties of the new retinal data output provided by the UGR group. Considered what aspects of the output were most important for integration with a cortical model. Explored literature for more advanced retinal models.	EPFL	Greg Francis (EPFL), Marc-Oliver Gewaltig (EPFL)	

Date	Description	Location	Participants	Comments
22.01.2014	Weber law update (SP11) Discussed plans for parallel paths of investigation. Greg will contact Eduardo about developing and benchmarking two (simple/complex) retinal models. The simple model will be used to investigate Weber's law for line lengths. Marc-Oliver will identify a master's student to start putting together a cortical model with retinotopic structure to receive visual inputs.	EPFL	Greg Francis (EPFL), Marc-Oliver Gewaltig (EPFL)	
26.05.2014	SP11 committee meeting	VC		

5.2 Scientific Coordination and Support: HBP Meetings

This table lists meetings between this SP and other SPs.

Date	Description	Location	Participants	Comments
10.04.2014	HBP SP9/SP11.3 In-Person-planning Meeting in Sabanci (Turkey)	Sabanci (Turkey)		
06.05.2014	SP9/SP11.3 JourFixe TelCo	VC		
03.06.2014	SP9/SP11.3 JourFixe TelCo	VC		
17.07.2014	HBP cross SP data, models, methods, and service exchange meeting (SP1, SP2, P3, SP4, SP5, SP6, SP7, SP8, SP9, SP10, SP11) Video conference for all SPs regarding their data and data-service offers and needs.	VC		
24.07.2014	SP9/SP11.3 in-person-meeting. Full day quarterly physical meeting of the SP9 and SP11.3 participants in Berlin.	Berlin		
05.08.2014	SP9/SP11.3 JourFixe TelCo	VC		
29.09.2014	HBP Summit 2014 Physical Meeting SP9-WP11.3	VC		

5.3 Scientific Coordination and Support: Monitoring & Quality Control

SP11 is coordinated by Karlheinz Meier (UHEI), who represents SP11 on the HBP Board of Directors. His team coordinates and monitors the development of the Applications Subproject and makes regular progress reports to the central HBP management.

The integration of the new Partners was a major goal in the second half of the period and this process started with a videoconference with all SP11 Partners in 05/2014. The SP manager informed the new Partners about the general purpose of this Project, it's relation to HBP and about reporting duties, namely the quarterly reports and the next Deliverables. The last quarter of the period was dominated by the preparation of the Deliverables D11.4.1 and D11.4.2 and the definition of Functions and KPIs. The Task leaders now regularly report the values via a central website to the Science & Technology Office at UHEI.

Communication among SP11 Partners is facilitated by an internal mailing list.

Annex A: Milestones

No.	Milestone Name	WP	Month Due	Month Achieved	See Page
MS200	Preparation of simplified virtual sensor and motor models for virtual robots	11.1	6	6	
MS208	Specification of benchmark tasks	11.3	6	6	
MS201	Experimental design for the first experimental task	11.1	12	12	
MS204	Standardised description format for biological signatures of brain disease	11.2	12	12	
MS209	Evaluation of suitable network architectures for benchmark tasks	11.3	12	15 est.	
MS310	Initial sensor interfaces operative	11.3	12	12	
MS315	Acquisition and interface Platforms and database	11.3	12	19	
MS308	CABot3 on SpiNNaker	11.3	15	15	

Annex B: Scientific Key Performance Indicators (SKPIs)

SP11's KPI information can be viewed at the Science and Technology Office's (STO's) KPI website via the link below:

https://flagship.kip.uni-heidelberg.de/jss/CollectKPI?uI=268&s=UJuR3AgTezrb&um=sPO&oSP=11

Annex C: References

[1] Data Mining- Concepts and techniques, Jiawei Han and Micheline Kamber, Elsevier 2006

[2] J. A. Pérez-Carrasco, B. Zhao, C. Serrano, B. Acha, T. Serrano-Gotarredona, S. Chen and B. Linares-Barranco, "Mapping from Frame-Driven to Frame-Free Event-Driven Vision Systems by Low-Rate Rate-Coding and Coincidence Processing. Application to Feed-Forward ConvNets," IEEE Trans. on Pattern Analysis and Machine Intelligence, vol. 35, No. 11, pp. 2706-2719, Nov. 2013

[3] Braitenberg, V.; Schüz, A.: Anatomy of the cortex. Statistics and geometry. Springer-Verlag, Berlin, 1991.

[4] Braitenberg, V.: Cell assemblies in the cerebral cortex. In R. Heim and G. Palm, editors, Lecture notes in biomathematics (21). Theoretical approaches to complex systems, pp. 171–188. Springer-Verlag, Berlin Heidelberg New York, 1978.

[5] Schmidt, M., Rückert, U.: Content-Based Information Retrieval using an embedded Neural Associative Memory. 9th Euromicro Workshop on Parallel and Distributed Processing Mantova, Italien, pp. 443-450, 2001.

[6] Steinbuch, K.: Die Lernmatrix. Kybernetik, 1:36–45, 1961.

[7] Willshaw, D.J., Buneman, O.P., Longuet-Higgins, H.C.: Non-holographic associative memory. Nature, 222, pp. 960–962, 1969.

[8] Palm, G.: On associative memories. Biological Cybernetics, 36:19–31, 1980.

[9] Brüderle, D., et al.: Establishing a novel modelling tool: a python-based interface for a neuromorphic hardware system. Frontiers in Neuroinformatics, Vol. 3, Article 17, pp. 1-10, 2009.

[10] Knoblauch, A., Palm, G.: Pattern separation and synchronization in spiking associative memories and visual areas. Neural Networks, 14, pp. 763–780, 2001.

[11] Knoblauch, A., Palm, G., Sommer, F. T.: Memory capacities for synaptic and structural plasticity. Neural Computation, 22, 289–341, 2010.

[12] Ryad Benosman, Sio-Hoi Ieng, Charles Clercq, Ciara Bartolozzi, Mandyam Srinivasan, (2012) Asynchronous frameless event-based optical flow, Neural Networks. Volume 27, March, Pages 32–37

[13] C. Posch, D. Matolin, and R. Wohlgenannt, "A qvga 143 db dynamic range frame-free pwm image sensor with lossless pixel-level video compression and time-domain cds," Solid-State Circuits, IEEE Journal of, vol. 46, no. 1, pp. 259 –275, jan. 2011

[14] Henri Lorach, Serge Picaud, Sio-Hoi Ieng, Jose Sahel, Ryad Benosman. Artificial retina : the multichannel processing and temporal dynamic of the biological system achieved with an asynchronous model, Journal of Neural Engineering (2012).

[15] Rogister et al., 2012) Rogister, P., Benosman, R., Ieng, S. H., Lichtsteiner, P., & Delbruck, T. (2012). Asynchronous event-based binocular stereo matching. Neural Networks and Learning Systems, IEEE Transactions on, 23(2), 347-353.

[16] João Carneiro, Sio-Hoi Ieng, Christoph Posch, Ryad Benosman, "Asynchronous Event-Based 3D Reconstruction From Neuromorphic Retinas". Neural Networks, 15th march 2013. DOI:10.1016/j.neunet.2013.03.006