Project Title:	**The Human Brain Project**
Project Acronym:	**HBP**
Type of instrument:	**FET Flagship CP/CSA**
Grant Agreement number:	**604102**

Response to the HBP Competitive Call for additional beneficiaries

Call Identifier: HBP Competitive Call for additional beneficiaries

Call topic 6 Neural configurations for neuromorphic computing systems

Neuromorphic Embodied Agents that Learn

NEAL

PROPOSAL PART B

Version: Oct. 18

Executive Summary

The goals of this work are: 1: place an embodied agent driven by spiking neurons on two neuromorphic platforms, SpiNNaker and HICANN. 2: extend the agent to learn categories of items from its environment that will improve its performance and communication with a user; and 3: do this in a neuro-psychologically plausible manner. Placing the agent on the chips will be a relatively straight-forward engineering task. Extending the agents learning abilities is linked to neuro-psychological plausibility. Exploration of learning over large numbers of neurons will involve a complex search of a vast problem space of topologies, learning algorithms, and neural models. The system will learn cell assemblies (CAs), the neural basis of categories, that are both effective and more closely aligned to psychological behaviour.

This will be done in two main phases, both involving software development, and handshaking with the neuromorphic platforms. The first phase will involve initial translation of an existing neuronal agent to PyNN along with resolving any inconsistencies with the hardware. The Cell Assembly robot version 3 (CABot3) will be translated to PyNN, SpiNNaker and HICANN, and runs in a virtual environment. There will be handshaking between the 3D environment and the neuromorphic chips. This first phase will result in CABot3 on both chips. The second phase will then exploit the hardware for exploring learning CAs and plans from the environment. CAs will be learned for both categories and instances of objects in the environment. These will fire persistently for times consistent with short-term memory (STM), and will help the agent in performing its tasks. This experimental search will include, but not be limited to, a range of Hebbian learning mechanisms, topologies, and neuron types. Simulations, both hardware and software based, will be used to find parameters so that the system can learn to perform correctly. Analysis of the topology and the evolution of the topology will drive the search. In addition to performance in the environment, this will involve exploration of STM dynamics and classification. Both categorisation and STM will be evaluated as cognitive models.

This project, NEAL, is an application on the neuromorphic platform; it can be extended to provide new agents, and new models of cognitive domains and brain areas. So, it will be one software tool from the brain simulation platform, a brain simulation engine that can run on neuromorphic chips. It will eventually lead to a contribution on the neuro-robotics platform, as extensions will be validated brain models embodied in physical robots.

Using a point model, NEAL will simulate billions of neurons in real-time on SpiNNaker, and the explorations of multi-area learning that persists for days will be an application of the brain simulation platform. Similarly, this project will be applicable to neuro-robotics in the cognitive areas that are being developed, and in the future in actual motion. This will integrate with the cognitive architectures, and the mathematical and theoretical foundations of brain research sub projects. The overall CABot network is designed by partitioning into sub-networks, and these can be grouped together to allow domains, and thus larger areas, to collaborate. In combination with the CA work, there is a strong link to cognitive architectures. Understanding the long-term dynamics of CA formation, and the short-term dynamics of CA activation and competition will advance the theoretical foundations. When possible, and for inspiration, data from the strategic human and mouse brain data sub projects will be used for, for example, topology, neural models and learning rules. We are particularly interested in data involving learning concepts and neural and synaptic changes. Though the project is part of the neuromorphic platform and sub project, it is also linked with the brain simulation and neuro-robotics platform, and the cognitive architectures, theoretical, and mouse and human brain data sub projects.

The project will contribute to the HBP by providing an extensible agent that can be used in a 3D environment, and by providing advancements for the HBP's neuromorphic systems to learn CAs. 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 neural language, vision, planning, and action modules, which have a reasonable degree of modularity. Improved models of CA learning will provide insight into the theoretical problem of concept formation in neural systems. This will be linked to the environment, psychology and known and posited neural behaviour, leading to a significant impact in both the short and longer term.

TABLE OF CONTENTS

B0. Cost and funding breakdown

Organisation 1 - Organisation Name: Middlesex University

Categories	**RTD (€)**
1. Personnel costs	90300
2. Travel & subsistence	6500
3. Durable equipment	0
4. Consumables	0
5. Subcontracting (detailed justification is needed)	0
6. Other direct costs	0
7. Total direct costs (Sum of row 1 to 6)	96800
8. Indirect costs (according to the organisation's cost model)	58080
9. Total costs (Sum of row 7 and 8)	154880
10. Requested EC contribution	116160

B1. Proposed Plan

B1.1. Objectives and approach

Neuromorphic Computing Systems Objective:	**Achieved by:**
Implement novel computing paradigms	Implementing an embodied agent in spiking neurons.
Generic circuit concepts of spiking neurons	Develop improved cell assembly models.

Operational Objective Area:	**Achieved by:**
Neuromorphic platform	Implementing simplified versions of brain models.
Theoretical foundations	Furthering generic CA models.
Brain simulation platform	Building point models for simulating brain areas.
Cognitive architecture	Building neuro-cognitive models to extract principles.

NEAL will implement a novel computing paradigm, an embodied agent in spiking neurons on neuromorphic platforms. The earliest phase will make use of the neuromorphic platform to implement a simplified version of a brain model by translating the existing CABot3 system to PyNN and then onto the neuromorphic chips. This relies on the existing CABot3 code, written in Java, but also on PyNN, and eventually on SpiNNaker and HICANN. CABot3 is implemented entirely in our own Fatiguing Leaky Integrate and Fire (FLIF) neural model written in Java. It uses over 100,000 neurons but simulates the entire system in roughly real-time on a PC since the discrete time steps correlate with 10ms of time. CABot3 is broken into 46 subnets, enabling some degree of modularity. We have already implemented the model without fatigue in PyNN. As several of the subnets do not take advantage of fatigue, several subsystems (e.g. the planning subnets) can begin to be implemented immediately. We expect that the fatigue model will be relatively easy to implement in standard PyNN neural models. Moreover, as part of CA learning a different fatigue model will be used, so this will need to be implemented. Learning is a similar problem because CABot3 uses both long and short-term synaptic modification (Huyck, 2009). We expect long-term learning to be relatively straight-forward, but short-term plasticity may be more of a challenge. Moreover these challenges may need a different solution for each of the chips. These proposed initial experiments will use up to 100,000 neurons but this should not break the “maximum of a few tens of thousands of neurons” mentioned in the call, because the time constant is so large. Similarly, the proposed system does not take advantage of STDP; learning is based entirely on co-firing in a given 10ms cycle.

The second phase will explore CAs, a generic circuit concept. This will meet the theoretical foundations objective by furthering generic CA models. During the development of CABot3, we showed that a large enough spiking neural network is Turing complete (Byrne & Huyck, 2010), and implemented an embodied agent solely in spiking neurons. During the development, it became clear that the difficult question was not how to program a neural system to function, but which neural model, network and topology were needed so that it could learn to function. Prior to CABot and since finishing, we have spent some time studying learning with these systems. Our earlier learning work (Huyck & Orengo, 2005; Huyck, 2007; Nadh & Huyck, 2012) focused on nets where all neurons were stimulated by the environment. The extension of our FLIF model to allow spontaneous firing from hypo-fatigue has enabled the growth of neural circuits into areas that are not directly stimulated from the environment (Huyck & Mitchell, 2013). This has still, in essence, been a sensory task since firing stops when the stimulus stops. CABot3 needed CAs to persist mainly for relatively precise times to support natural language parsing dynamics, but also for planning. CAs are a long standing and well supported theory linking psychology and biology (Hebb, 1949; Huyck & Passmore, 2013). In CABot3, CAs were mostly programmed by setting synaptic weights. Some progress has been made by using small world topologies so that the CAs persistently fire longer when they are activated more strongly, and for longer if reactivated. NEAL will combine these two strands so that CAs, that persist for psychologically realistic times, can be learned.

Some of the evaluations of NEAL will meet the cognitive architecture objective by building neuro-cognitive models to extract generic principles. CABot3 already has a cognitive model of parsing and one of rule learning. The new simulations will learn novel categories of visual items from the environment. There will be a range of 3D shapes with a range of visual textures. These will initially be linked to labels, but later learned in an entirely unsupervised fashion. Categories will have a special behaviour in the environment. We have a neuro-cognitive model of a rule learning system that makes use of reinforcement learning from environmental feedback (Belavkin & Huyck, 2010), and this mechanism will be used to determine the categories. A classic experiment (Shepard, Hovland, & Jenkins, 1961) will be used as a cognitive test. CABot3 plans using a Maes net (Maes, 1989) that has been programmed (in FLIF neurons), but NEAL will learn new plans. Initially this will be by integrating our reinforcement mechanism with the Maes net, and by direct user instruction. The existing net is built of orthogonal CAs (sharing no neurons), and the reinforcement mechanism will modify the strength of associations between CAs. The user instructions will create new CAs. This limited approach will be replaced by an overlapping model where basic units will be learned associating goals, actions, and facts from the environment. This will include extra subnets to generate and evaluate plans supporting the generation of more complex plans. NEAL will also take advantage of the larger number of neurons. We are looking for long term systems lasting days, learning throughout. This is a developmental neuro-psychology problem. A comparison of the two systems using STDP combined with a compensatory mechanism for the CA formation problem is particularly promising. The variability of analogue neurons and synapses may lead to particularly powerful attractors, and CAs are attractors.

The system meets an objective of the brain simulation platform; it is a brain model based on a point model level of description. CA learning really is open ended research. We will develop a working system that generalises, explore the space of working systems, and work in promising areas. Exploiting the variability of neurons and synapses both analogue and discrete, we will explore ranges of options. There are a huge number of options so we will be guided by the neuroscience as much as possible, but meet the milestones. This relies only on the earlier systems. We will interact with the HBP community to learn what is neuro-psychologically plausible, and what can be implemented.

B1.2. Progress beyond the state of the art

Learning CAs, and learning plans are advances on the state of the art. However, translating the FLIF model, and translating CABot3 are not really beyond the state of the art. Instead these two tasks will focus on aligning the current state of the art CABot3 system with the HBP. The simple FLIF model that we have developed will be readily comparable with other neural models both in software and hardware simulations. CABot3 will be connected via PyNN to neuromorphic chips that will support a system with a much larger number of neurons and synapses.

Learning CAs is novel not because, like earlier work, these CAs will be linked to the environment, but they are novel because of the psychological nature of these CAs and their effectiveness. These CAs will persist for times similar to short-term psychological memory. An initial metric here is the ACT-R times (Anderson & Lebiere, 1998), but we will use further evidence. These CAs will be derived from the environment, and help the agent improve its performance in that environment This will be a step toward a learning system capable of learning domain specific and general concepts that have rich semantics.

Learning plans is also novel not because systems cannot learn plans, but because existing neural systems that learn plans are limited. This will provide novel mechanisms for learning plans in a neural system, and a novel means of learning CAs, as they are a core component of our Maes net plans.

For example, the domain for the second agent could be a mushroom game, where the agent moves about the 3D virtual environment collecting mushrooms. It will need food and sleep, will get injured, and be able to attack and collaborate with other agents in the environment. The mushrooms will come in varying shapes, colours and textures, and the agent will classify them, and discover their properties. It will also learn better plans for getting the appropriate mushrooms, and for interacting with other agents.

Finally, the agents provide one possible framework for building integrated neuro-cognitive agents. The NL parsing component of CABot3 was a neuro-cognitive model of parsing (Huyck, 2009). It made use of the tripartite theory (Jackendoff, 2002), a linguistic theory based on the lexical, syntactic and semantic domains. The theory says that there are core processes, and processes for communicating between domains. These also apply to other domains (e.g. morphology), but it seems reasonable to extend the theory to areas outside of language processing. Our development methodology breaks the overall net into subnets. In some cases, though currently only a few and not much more than metaphorically, these subnets can be linked directly to brain areas. Additionally, subnets provide a plug-and-play mechanism. One researcher can develop one subsystem (e.g. vision), and as long as it communicates with spikes, it can be replaced by another version relatively easily.

B1.3. Methodology and associated work plan

The overall strategy is to build working systems. Initially we will build agents functioning in an environment, task 1.1 (T1.1) and T1.2. Later we will expand these agents (T1.3 and T1.4) so that they can learn more about their environment, and thus be more useful in that environment. While doing this, working agents will perform tasks as cognitive models; similarly; this will have some correlation with biological data, though anything that is particularly solid may be beyond the scope of this project.

We have discussed SpiNNaker with Furber for several years now, and have planned on putting our model onto it from our first conversation. Consequently, ideas of working on SpiNNaker are more fully formed than those for working on HICANN, and the plan reflects this bias.

The NEAL project consists of one work package. That is broken into 4 Tasks (see table 1.2a below).

Task 1.1: Transfer FLIF model to PyNN and neuromorphic chips

We have already translated the Leaky Integrate and Fire component of our neural model to PyNN, and will immediately begin translating some of the CABot3 components (e.g. early vision and planning) to PyNN. Later, we will need to add the fatigue component to the model. A standard PyNN model should be sufficient. There are two variants of fatigue, and only the first, milestone 1.1 (M1.1), needs to be implemented for CABot3. The second will be translated after month 9.

Learning will be included in the system. PyNN has good support for LTP, and it will be included in the first few months. A critical component of the CABot3 model is binding via STP (M1.2). This should work in PyNN, and all basic neural and synaptic components will run on both chips (M1.3).

Finally, CABot3 is precisely timed and it is not clear how readily this will translate to either chip. We are hopeful that the 10ms integration constant will solve the problems, but may need to increase the number of neurons for CABot3. Even if CABot3 does not require more neurons, there is scope for exploration of the dynamics and robustness of neural processing circuits with more neurons. That is, by using more neurons, the systems will be more effective, and cope with more hardware failures.

Task 1.2: Transfer CABot3

CABot3 runs in a virtual 3D environment. CABot3 has a mobile avatar in the environment. Input from the environment comes in the form of a pixel image from the avatar's camera and text commands issued from a user. We will begin by translating the vision, control, and planning subnets of CABot3 into PyNN. These are not dependent on real-time behaviour, nor do they particularly need the fatigue model. Via Python, we will integrate the environment with the agent providing both input to and symbolic motion output from the PyNN based CABot3 (M2.1).

As the fatigue, LTP and STP primitives are implemented (T1.1), these will be integrated into the PyNN agent, then to the SpiNNaker agent (M2.2), and then the HICANN agent. All three agents will be tested and compared. Tests will include the ability to parse commands, view the environment, build a simple spatial cognitive map of the environment, and to learn which rule is a correct rule. These will be compared with the existing Java based system and each other. We expect that all will behave almost identically, though extra neurons could be used by the neuromorphic agents.

Task 1.3: Learn CAs

We will use the PyNN model to learn simple categories, then use the modified fatigue model to include spontaneous activation from the new fatigue model. This will bring us to the current state of the Java model on PyNN and the chips. We will then add new subnets replacing the higher level vision subnets, and the system will learn the visual categories by exploring. Instances of categories will be co-presented with labels so the system can learn words. Labelling provides a ready test, and links to the language system.

We will then develop a system that learns CAs that persist for the time psychological short-term memories persist (M3.1). If there is more evidence for a CA, it will persist longer. If a CA is reactivated, by for instance being presented again, it will persist longer the second time. Our current idea is that our existing sensory like learning mechanism can be integrated with subnets that support persistence and top-down effects. This will be compared to the ACT-R memory model.

A new environment will then be used; using video game technology it is quite easy to develop new 3D environments. This will be a more sophisticated task, searching for particular types of objects to fulfil current goals or needs. Environmental feedback will enable the system to learn things that will help its performance. We will test this on categories with varying degree of feature overlap (Shepard, Hovland, & Jenkins, 1961) to allow a cognitive test (M3.2). We will also test that its game performance improves.

Task 1.4: Learn Plans

We will change associations within Maes net elements to improve an existing plan driven by reinforcement learning. This will involve a modification of associations between existing CAs, which implement the Maes net. Similarly, new elements within the net will be created and linked in response to user commands, and the associations between these will be modified to improve the overall plan. Together these are M4.1.

After this, we will learn overlapping CAs for plans. The earlier plan, including the learned plans will use CAs that do not share neurons. The new plan will take advantage of shared neurons to learn more sophisticated, and neuro-psychologically plausible plans. This will be driven by separate plan generation subnets, and we will consider turning learning on and off by neuromodulators. This is a speculative venture, so it does not have a milestone. The system will be tested on the new 3D environment, which is not a cognitive test (M3.3 and M4.2).

Milestone Descriptions (see table 1.2b below)

Milestones are directly linked to tasks, and proceed sequentially through the task. T1.1 has three milestones. M1.1 is to put the CABot3 FLIF model onto PyNN by month 2. M1.2 is to include the STP model in PyNN by month 6. M1.3 is to have both FLIF models, and long and short-term synaptic modification on PyNN and both chips. M1.1 and M1.2 are necessary to move the project forward; M1.3 is more difficult but is also largely non-blocking.

There are only two milestones for the CABot3 task T1.2. M2.1 puts a simple version of the agent into PyNN by month 3; this will support movement onto SpiNNaker and HICANN. M2.2 has the complete agent on SpiNNaker via PyNN; it takes advantage of the first task. CABot3 will then be implemented on HICANN.

The largest task is T1.3, learning CAs. This is much more exploratory than T1.1 and T1.2, but at least these three milestones will be met. M3.1 is properly persistent learned categories. M3.2 is a cognitive model of a classification task, and M3.3 improves the overall agent's performance. We are confident in a reasonable advancement with M3.3, but hope for a significant one.

Task T4 is learning plans, and thus is learning process. M4.1 is relatively straight-forward porting our existing reinforcement learning to the Maes net, and responding to user commands to add new elements to that net. M4.2 is a catch all for the project, allowing the project to wrap up with one agent, or a variant for each chip; this could include the more sophisticated planning system.

Table 1.2a: WP and Tasks description

Work package number:	WP 1	**Start month:**	M 1	**End month:**	M 24
Work package title:	Embodied Agent Development and Evaluation				
Activity type:	**RTD**				
Participant Number:	*1 Middlesex University*				
Participant Short Name:	MU				

Objectives:

1: Provide an embodied cognitive agent in spiking neurons on SpiNNaker and HICANN that can be readily modified and extended.

2: Extend the agent to learn environmentally useful and neuro-psychologically realistic cell assemblies.

Description of Work and role of the partners:

Task 1.1: Transfer FLIF model to PyNN and neuromorphic chips (months 1-12)

All will be executed by Middlesex (MU). This will take 4 person months, and use Huyck and the RA.

Translate both variants of the fatiguing FLIF model to PyNN and SpiNNaker and HICANN. Integrate LTP, and STP models with PyNN and the chips.

Task 1.2: Transfer CABot3 (months 1-9)

MU: This will take 4 person months, and use the RA, Mitchell, and Huyck.

Take existing CABot3 Java code and move portions based on LIF model to PyNN. Integrate the 3D environment with PyNN and the chips. Translate the full CABot3 system to PyNN and chips. As soon as the chips become available, the PyNN agent will be integrated with SpiNNaker, and then HICANN.

Task 1.3: Learn CAs (months 6-24)

MU: The most time will be spent on this task taking 7 person months, by Huyck and the RA.

Explore the short and long term dynamics of CA persistence and creation. Use compensatory Hebbian learning with subnets of differing topologies. Explore varying neural models and STDP. A new 3D environment and task will support the development of a more sophisticated agent that can learn a wider range of semantics.

Task 1.4: Learn Plans (months 13-24)

MU: This will take 4 person months by Mitchell, the RA, and Huyck.

Modify existing plans in response to environmental feedback. Create new plans and plan elements from user instruction. Expand plan capability by learning overlapping CAs for plan elements.

Detailed allocation of effort (person months)

Proposer Number	Proposer Name	Person Months
1	MU	19

Table 1.2b: Milestones description

Milestones number	Milestones name	Lead proposer short name and number	Delivery month	Comments
1.1	FLIF in PyNN	MU 1	2	
1.2	STP in PyNN	MU 1	6	
1.3	Model on Chips	MU 1	12	Both Chips
2.1	Simple Agent on PyNN	MU 1	3	
2.2	CABot3 on SpiNNaker	MU 1	9	HICANN date flexible
3.1	Learned CAs persist like STM	MU 1	15	Based on ACT-R model
3.2	Classification Cognitive Model	MU 1	18	Based on Shepard et al.
3.3	Learned CAs help the agent	MU 1	21	Game based evaluation
4.1	Weight plans and cache commands	MU 1	18	
4.2	Agent complete	MU 1	24	Possibly 1 version for each chip

B2. Implementation

B2.1. Participants

Organisation 1

Name of the Organisation / Department	Middlesex University Higher Education Corporation (Department of Computer Science)
Location	London, UK
Description of the Organisation / Department (300 words limit)	Middlesex University is a thriving British University based in London with a large international network. The University teaches 40,000 students. The project team is based in the Department of Computer Science within the School of Science and Technology. Middlesex University, the School, and the Department have extensive experience of European funding including Framework Programmes, and non-European funding. Supported by the Research and Knowledge Transfer Office; MU has extensive expertise as a coordinator of FP projects. The Department is currently coordinating two FP7 projects, 'CRitical Incident management training System using an Interactive Simulation environment (CRISIS)' (contract no.: 242474) and 'Warehousing images in the digital hospital: interpretation, infrastructure, and integration (WIDTH)' (contract no.: 269124). It has recently been given a funded FP7 collaborative project VALCRI - Visual Analytics for Sense-making in Criminal Intelligence Analysis - worth €13 million to the consortium. The Department was a partner in 'Emerging Technoethics of Human Interaction with Communication, Bionic, and robotic Systems (ETHICBOTS)'(contract no.: 017759), In addition to these projects, the Department and School have participated in a range of grants and research that is particularly relevant to the NEAL project. There is a large Artificial Intelligence research group, and a large robotics group. The AI group has several members active in neuronal systems including: Professor Andreas Albrecht is working in attractor dynamics, Dr Belavkin in neuro-cognitive systems, Dr Kahn in compartmental neural models, Dr Passmore in cell assemblies and virtual environments, and Dr Yang in spiking models for central pattern generators. The robotics group includes several researchers interested in neuro-robotics including Dr Yang, and Professor Martin Smith. The school also includes many PhD whose research could be used in this project. Similarly, MSc and BSc thesis students could work on thesis projects directly linked to the NEAL project.

Previous experience relevant to the tasks the participant will undertake in the project.	The MU project team has extensive experience developing large software packages. In particular, the development of CABot3 demonstrates our ability to deliver large spiking neural network systems. Huyck has also worked on GATE, the General Architecture for Text Engineering. While at Microsoft, Huyck developed Visual Basic. He was invited to speak at the Telluride Neuromorphic Engineering Workshop in 2011. While there, there was some exploration of SpiNNaker's 4 chip model. Huyck has managed several grants including the UK EPSRC CABot3 grant 'Natural Language Parsing with Cell Assemblies: computational linguistics with attractor nets' and 'Modelling Cell Assemblies as a Neuro-Psychological Phenomenon and for Practical Applications'.

Title	Professor of Artificial Intelligence
First Name	Christian
Last Name	Huyck
Profile (300 words limit)	Huyck thinks we are a long way from a full-fledged Turing test passing AI, but that the best way to get there is to follow the human model, both psychologically and neurally. Huyck has been pursuing that path for 20 years now (Huyck, 2001), and has made some strides. He thinks it rather unlikely that he will complete this goal by himself, or even leading a small group. The HBP, however, fits into this plan perfectly. Huyck completed his PhD at the University of Michigan in 1994 with a thesis on human like natural language parsing. Having participated in the Message Understanding Competitions, Huyck continues research in Computational Linguistics. While attempting to resolve the Prepositional Phrase Attachment Ambiguity problem, he realised semantics was needed; he knew about CAs and that people use them to for their own semantics to resolve attachments. Naively, he thought he would just put in existing CA algorithms, and the problem would be solved. Fifteen years later, this problem was solved with CAs (Nadh & Huyck, 2012). The CA based solution is currently the best algorithm for resolving this ambiguity. Initially, Huyck was surprised that these algorithms did not already exist. Seeing that CAs needed to be understood, Huyck set about developing a neural model and learning algorithms. Obvious problems with full-fledged neural systems, such as the binding problem, have been addressed. While developing CABot agents, he realised that CA systems were Turing complete, and proved it.

	Huyck has been at Middlesex University since 1998, first as a Senior Lecturer, then Principal Lecturer, then Reader, and is currently Professor of Artificial Intelligence. He formed the AI research group on arriving, and the group has continued to grow throughout his tenure. He continues to pursue his goal of developing a real AI by following the human model.

Title	Dr
First Name	Ian
Last Name	Mitchell
Profile (300 words limit)	Mitchell has interests in Biologically Inspired Algorithms, artificial neural networks (ANNs), and combinations of the two. Mitchell completed his PhD at University of North London in 1999 on sequence recognition using a novel ANN (Bavan & Mitchell, 2000), known as Graph-Set And Associative Memory, GSAAM. GSAAM based the structure of the network on the sequences in the training set data. The weights were decided by a modified back-propagation algorithm. After training, recall would involve presenting an unseen partially incomplete sequence to GSAAM, which then recalled multiple candidate solutions each with an associated probability. Mitchell went on to compare ANNs and other techniques used in categorisation (Cairns, Huyck, Mitchell, & Wu, 2001). Categorisation continues to be an interest, with recent work comparing Kohonen nets with biologically realistic nets on this task (Huyck & Mitchell, 2013). Mitchell then pursued a new direction towards Evolutionary Computation with the intention of merging this with ANNs. This started looking at new ways to represent problems using Genetic Algorithms (Mitchell & Pocknell, 2000; Agrawal, Mitchell, Passmore, & Litovski, 2005). Mitchell retained his interests in ANNs using them to categorise and predict sequences (Jiang, Wu, & Mitchell, 2006) and continues to explore ANNs (Huyck & Mitchell, 2013; Tian, Guo, Liu, Mitchell, Cheng, & Zhao, 2013). Mitchell has been at Middlesex University since 1998, first as a Lecturer, then Senior Lecturer and is currently a Principal Lecturer. Mitchell and Huyck formed the AI research group on their arrival at Middlesex. It now contains 12 faculty, and roughly the same number of PhD students. Mitchell has been Director of Programmes since 2007 and introduced and developed a number of programmes, including Computer Forensics. In August 2013 Mitchell relinquished his role of Director of Programmes with the intention to return to research in AI-related fields.

B2.2. Resources to be committed

Huyck will be funded for 4.8 person months, Mitchell for 2.4 person months, and the RA (Wieclaw) for 12 person months. Huyck will manage the project, participate in all 4 tasks, and lead T1.1, T1.2 and T1.3. Mitchell will lead T1.4, and the planning portion of T1.2. The RA will focus on T1.2 then T1.3, but spend some time on T1.1 and T1.4.

This neuromorphic work complements Huyck's main research stream, so extra time could be devoted to the project beyond the budgeted time. Similarly, Wieclaw's doctoral work complements the project, so additional time could be devoted beyond the budget. Additionally, other staff at Middlesex (e.g. Dr Zhijun Yang, Dr Roman Belavkin, Dr Peter Passmore, and Dr Nawaz Khan) are involved in closely related work, so could contribute to the project in an unbudgeted fashion.

The overall project is relatively flexible with relatively few dependencies. The tasks are relatively separated from each other, though the later parts of T1.2 do depend on parts of task T1.1, as does T1.3. The separate subnets and domains enable team members to move the project forward independently. Huyck is driving most of the research, and is available to support Mitchell and Wieclaw. Similarly, Wieclaw does have experience with the FLIF model, but could be replaced by another relatively skilled developer if necessary.

The time budgeted for T1.1 and T1.2 should be sufficient in itself. T1.3 and T1.4 are more open ended. Development will proceed within the budgeted time and milestones will be met. Additional (unbudgeted) resources could allow the milestones to be met early, and for the results to be more significant.

Travel expenses of 6500 euros will be used for travel to conferences and other HBP sites.

B2.3. Feasibility of the work Plan within the available time frame

Huyck will manage the project, supporting Mitchell and the RA. It is expected that the half-time RA will be Anna Wieclaw, who is also working on classification with CAs. Huyck will actively communicate with the larger HBP, and encourage discussion between the entire NEAL team and the HBP community.

Initially, the team will work on tasks 1 and 2, resulting in CABot3 on PyNN and both chips. Having already translated the basic LIF model to PyNN, the project will be able to start on April 1st 2014. Initially, Huyck will concentrate on resolving fatigue and learning problems, while Mitchell and the RA will concentrate on translating CABot3, and integrating the environment. Regular meetings will support ad hoc meetings to move the project forward.

If the FLIF model including STP can be readily translated to PyNN and then onto the chips, T1.1 and T1.2 will go smoothly. There are two possible problems: the fatigue models and STP. Discussion with Furber indicate that SpiNNaker can implement both, though STP may be computationally expensive. The desired and expected course of action will be to translate the FLIF models to the adaptive exponential integrate and fire model on PyNN, and make use of existing PyNN STP processes that directly translate to HICANN (and SpiNNaker). However, this may not be feasible. It is possible to translate the FLIF model to a PyNN backend, but that backend would have to be resolved with HICANN. These problems may need to be resolved explicitly with the chip development teams; as the already implemented Java code is relatively simple, there is no theoretical problem, but there may be practical difficulties. Similarly, large portions of the system could be rewritten to work with a LIF neural model. If the FLIF model cannot be translated, resources will be switched from learning (T1.3 and T1.4) to agent development T1.2; if STP cannot be implemented, and thus binding cannot be easily implemented, the context free grammar parser will be replaced with a regular grammar parser. None the less, even in this worst case scenario, embodied agents will run on both chips.

Note that we feel the schedule for tasks T1.1 and T1.2 are generous. We are confident of making the milestones and fulfilling these objectives. Additional time made available will be used to explore the two remaining tasks, the options provided by the hardware (e.g. STDP), and the larger HBP.

Both Tasks T1.3 and T1.4 are about learning CAs. While we have made consistent advances in this area for quite some time, it is not clear how well things will proceed. We are confident that the milestones will be met, but there are a range of possible successful outcomes from learned CAs categorising shapes to CAs involved in plans that can learn plans in the environment that beat a human in a contest and even insightful plans.

Undone consider Rui thinks the advancement beyond state of art could have a figure.

B3. Expected Impact

There are three main ways that this project will increase the impact of the HBP: a working agent will support the cohesion of the project; new understanding of CAs will be gained; and more sophisticated and valuable AI systems will be developed. NEAL will have an impact by itself, but crucially it provides support for these further advancements.

Firstly, the development of an embodied cognitive agent in spiking neurons, running on neuromorphic chips will support the development of more sophisticated agents later in the HBP. Neural cognitive architectures, neural robotics, and the understanding of the brain all require huge numbers of neurons working together. Simple versions of many cognitive domains can be developed, potentially linked to brain areas, and NEAL provides several of these domains. Later in the HBP, other domains can be developed, more sophisticated domains can be developed, and different alternatives to domains can be developed. Huyck is particularly keen on building a language system that learns language and is linked to semantics. These can then be combined in a CABot-like system, and evaluated to move the HBP forward. While NEAL is based on a point neural model, it and its components can be integrated with more compartmental models, and other point models because communication is based on spikes. Neural cognitive models can be more closely linked to biological data. For this impact, these areas and agents will need to be developed.

Secondly, the impact of the HBP will be increased by extending the understanding of CAs. CAs are crucial neuro-psychological building blocks, each composed of many neurons. The scientific community does not fully understand how they are formed, how they act, or how they interact. NEAL will further this understanding, and provide the basis for a further growth in understanding. NEAL will further the understanding of short and long term CA dynamics, how they compete, and how the collaborate. Work beyond NEAL could extend this in a variety of ways including, for example, development of cognitive maps, integration with and relation to central pattern generators, and relations between semantic and episodic CAs. Each extension of understanding will increase the range of possible applications of these agents. The impact will be increased by discussion within the HBP, and links to biological and neuro-psychological data.

Thirdly, one impact of the HBP will be better computer systems. Sophisticated spiking neuron cognitive agents will be able to learn domains and avoid the brittleness of current AI systems. These agents will assist people as, for example, video game agents, improved interfaces, and improved data mining. For this to happen, cognitive agents will need to be developed during the HBP. Spiking neuron cognitive agents can then be used for the basis of industrially viable applications during and after the project. These can be based on neuromorphic hardware, supercomputers, distributed computation, and even with commonly available hardware including PCs, game consoles, or mobile phones. The FLIF model is computationally inexpensive, and understanding derived from the HBP could lead to robust systems that function on low-end hardware. At the other extreme, agents could integrate neuromorphic hardware, supercomputers and distributed computation to lead to full-fledged Turing test passing AI and maybe beyond.

This blue sky thinking will not overshadow our development of NEAL. The project meets the objectives of the neuromorphic additional beneficiaries call, and helps to meet several objectives of the larger HBP. It meets both of the implicit objectives of the call by being an implementation of a novel computing paradigm, an embodied agent in spiking neurons; and by exploring a generic circuit concept, CAs. It meets the strategic objective of the full flagship 1 (SOFF-1) by being an ICT platform for brain modelling that links a point neural model to neuromorphic hardware. It meets SOFF-2 by both modelling the brain over short and long term scales, and by creating brain-inspired technology. It meets SOFF-5, and to some degree SOFF-3, by using biological data to develop the agents, thus linking ICT with neuroscience. Finally, the NEAL team are committed to the collaborative nature of HBP and will contribute to SOFF-6 by collaborating with HBP partners whenever possible.

Bibliography

Agrawal, A., Mitchell, I., Passmore, P., & Litovski, I. (2005). Dynamics in Propotionate Selection. *Lecture Notes in Computer Science*.

Anderson, J., & Lebiere, C. (1998). *The Atomic Components of Thought.* Lawrence Erlbaum.

Bavan, S., & Mitchell, I. (2000). A Novel Connectionist Model for Pattern Directed Rule Based Programming. *Expert Systems*.

Belavkin, R., & Huyck, C. (2010). Conflict Resolution and Learning Probability Matching in a Neural Cell-Assembly Architecture. *Cognitive Systems Research*, 93-101.

Byrne, E., & Huyck, C. (2010). Processing with Cell Assemblies. *Neurocomputing*, 76-83.

Cairns, P., Huyck, C., Mitchell, I., & Wu, W. (2001). A Comparison of Categorisation Algorithms for Predicting Cellular Localisation Sites of Proteins. *Knowledge and Information Systems*, 296-300.

Hebb, D. (1949). *The Organization of Behavior.* New York: John Wiley and Sons.

Huyck, C. (2001). Cell Assemblies as an Intermediate Level Model of Cognition. In S. Wermter, J. Austin, & D. Willshaw, *Emerging Neural Architectures Based on Neuroscience* (pp. 383-397). Springer.

Huyck, C. (2007). Creating Hierarchical Categories Using Cell Assemblies. *Connection Science*, 1-24.

Huyck, C. (2009). A Psycholinguistic Model of Natural Language Parsing Implemented in Simulated Neurons. *Cognitive Neurodynamics*, 316-330.

Huyck, C. (2009). Variable Binding by Synaptic Weight Change. *Connection Science*, 327-357.

Huyck, C., & Mitchell, I. (2013). Compensatory Hebbian Learning for Categorisation in Simulated Biological Neural Nets. *Biologically Inspired Cognitive Architectures*, 3-8.

Huyck, C., & Orengo, V. (2005). Information Retrieval and Categorisation Using a Cell Assembly Network. *Neural Computing and Applications*, 282-289.

Huyck, C., & Passmore, P. (2013). A Review of Cell Assemblies. *Biological Cybernetics*, 263-288.

Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution* . Oxford University Press.

Jiang, W., Wu, X., & Mitchell, I. (2006). MESSM: A Framework for Protein Fold Recognition Using Neural Networks and Support Vector Machines. *Bioinformatics Research and Applications*.

Maes, P. (1989). How to Do the Right Thing. *Connection Science*, 291-323.

Mitchell, I., & Pocknell, P. (2000). A Temporal Representation for GA and TSP. *Lecture Notes in Computer Science*, 651-660.

Nadh, K., & Huyck, C. (2012). A Neurocomputational Approach to Prepositional Phrase Attachment Ambiguity Resolution. *Neural Computation*, 1906-1925.

Shepard, R., Hovland, C., & Jenkins, H. (1961). Learning and Memorization of Classifications. *Psychological Monographs*, 1-42.

Tian, K., Guo, B., Liu, G., Mitchell, I., Cheng, D., & Zhao, W. (2013). KCMAC-BYY: Kernel CMAC using Bayesian ying-Yang Learning. *Neurocomputing*, 24-31.

B4. Ethical issues

Describe any ethical issues that may arise in the action, filling the following form below:

	YES	NO	PAGE
Informed Consent			
• Does the proposal involve children?			
• Does the proposal involve patients or persons not able to give consent?			
• Does the proposal involve adult healthy volunteers?			
• Does the proposal involve Human Genetic Material?			
• Does the proposal involve Human biological samples?			
• Does the proposal involve Human data collection?			
Research on Human embryo/foetus			
• Does the proposal involve Human Embryos?			
• Does the proposal involve Human Foetal Tissue / Cells?			
• Does the proposal involve Human Embryonic Stem Cells?			
Privacy			
• Does the proposal involve processing of genetic information or personal data (e.g. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)			
• Does the proposal involve tracking the location or observation of people?			
Research on Animals			
• Does the proposal involve research on animals?			
• Are those animals transgenic small laboratory animals?			
• Are those animals transgenic farm animals?			
• Are those animals cloned farm animals?			
• Are those animals non-human primates?			
Research Involving Developing Countries			
• Use of local resources (genetic, animal, plant etc)			
• Impact on local community			
Dual Use			

• Research having direct military application			
• Research having the potential for terrorist abuse			
ICT Implants			
• Does the proposal involve clinical trials of ICT implants?			
I CONFIRM THAT NONE OF THE ABOVE ISSUES APPLY TO MY PROPOSAL	Yes		