



HBP
The Human Brain Project

The Human Brain Project

**A Report to
the European Commission**

Abstract

Understanding the human brain is one of the greatest challenges facing 21st century science. If we can rise to the challenge, we can gain fundamental insights into what it means to be human, develop new treatments for brain diseases and build revolutionary new Information and Communications Technologies (ICT). In this report, we argue that the convergence between ICT and biology has reached a point at which it can turn this dream into reality. It was this realisation that motivated the authors to launch the Human Brain Project – Preparatory Study (HBP-PS) – a one-year EU-funded Coordinating Action in which nearly three hundred experts in neuroscience, medicine and computing came together to develop a new “ICT-accelerated” vision for brain research and its applications. Here, we present the conclusions of our work.

We find that the major obstacle that hinders our understanding of the brain is the fragmentation of brain research and the data it produces. Our most urgent need is thus a concerted international effort that can integrate this data in a unified picture of the brain as a single multi-level system. To reach this goal, we propose to build on and transform emerging ICT technologies.

In neuroscience, neuroinformatics and brain simulation can collect and integrate our experimental data, identifying and filling gaps in our knowledge, prioritizing and enormously increasing the value we can extract from future experiments.

In medicine, medical informatics can identify biological signatures of brain disease, allowing diagnosis at an early stage, before the disease has done irreversible damage, and enabling personalised treatment, adapted to the needs of individual patients. Better diagnosis, combined with disease and drug simulation, can accelerate the discovery of new treatments, speeding up and drastically lowering the cost of drug discovery.

In computing, new techniques of interactive supercomputing, driven by the needs of brain simulation, can impact a vast range of industries, while devices and systems, modelled after the brain, can overcome fundamental limits on the energy-efficiency, reliability and programmability of current technologies, clearing the road for systems with brain-like intelligence.

The supercomputer and hardware technologies we need are rapidly improving their performance, following well-established industry roadmaps. In other essential technologies, European academic institutions already lead the world. The vision we propose would leverage these strengths, driving a

radical transformation of ICT and brain research, and enabling European bio-tech, pharmaceutical and computing companies to pioneer the development of what are likely to become some of the largest and most dynamic sectors of the world economy.

Realising this vision and ensuring a leading role for European companies and researchers will require a massive long-term research effort, going far beyond what can be achieved in a typical European research project or by any one country. As a foundation for this effort, we propose to build an integrated system of ICT-based research platforms, which without resolving all open problems, would allow neuroscientists, medical researchers and technology developers to dramatically accelerate the pace of their research.

Building and operating the platforms will require a clear vision, strong, flexible leadership, long-term investment in research and engineering, and a strategy that leverages the diversity and strength of European research. It will also require continuous dialogue with civil society, creating consensus and ensuring the project has a strong grounding in ethical standards.

We estimate that the total cost would amount to Eur 1,190 million, spread over a period of ten years. Of this sum, Eur 643 million would come from the European Commission, the remainder from other sources. We expect that, as the value of the platforms becomes apparent, this initial investment would trigger an avalanche of additional public and industrial funding, making it possible to perform research that goes beyond the precise limits we have indicated in this report.

These requirements can only be met by a project on the scale of a FET Flagship. In this report, therefore, we propose that the European Commission launches such a Flagship. We call it *The Human Brain Project* (HBP). We summarise the goal of the project as follows.

The Human Brain Project should lay the technical foundations for a new model of ICT-based brain research, driving integration between data and knowledge from different disciplines, and catalysing a community effort to achieve a new understanding of the brain, new treatments for brain disease and new brain-like computing technologies.

The Human Brain Project – Preparatory Study:

Partners: Ecole Polytechnique Fédérale de Lausanne (CH), Ruprecht-Karls-Universität Heidelberg (DE), Forschungszentrum Jülich GmbH (DE), Centre Hospitalier Universitaire Vaudois (CH), Karolinska Institutet (SE), Universidad Politecnica de Madrid (ES), Wellcome Trust Sanger Institute, Genome Research Limited – Genes to Cognition (UK), Technische Universität München - Fortiss GmbH (DE), Interuniversity Microelectronics Centre (IMEC) (BE), Hebrew University of Jerusalem (IL), Institut Pasteur (FR), Innsbruck Medical University (AT), Commissariat à l'énergie atomique et aux énergies alternatives (FR)

Coordinator: Henry Markram

Co-directors: Henry Markram, Karlheinz Meier

Executive Directors: Richard Frackowiak, Karlheinz Meier, Thomas Lippert, Sten Grillner, Henry Markram

External Science Advisory Board: Torsten Wiesel (Chair), Yves Agid, Andreas von Bechtolsheim, Bob Bishop, Sydney Brenner, André Hoffmann, André Syrota

Internal Advisory Board: Wolf Singer (Chair), Yadin Dudai, Thomas Schulthess

Division leaders: Anastasia Ailamaki, Katrin Amunts, Jean-Pierre Changeux, Javier DeFelipe, Stanislas Dehaene, Alain Destexhe, Kathinka Evers, Richard Frackowiak, Steve Furber, Wulfram Gerstner, Seth Grant, Sten Grillner, Jeanette Hellgren-Kotaleski, Alois Knoll, Thomas Lippert, Henry Markram, Karlheinz Meier

Editor: Richard Walker

Contributing Authors: Ailamaki A., Alvandpour A., Amunts K., Andreoni W., Ashburner J., Axer M., Baaden M., Badia R., Bartolome J., Baum L., Baumans V., Baumela L., Bernèche S., Biddiscombe J., Bjaalie J., Blanke O., Bogorodzki P., Bolten M., Born J., Brice A., Brunel N., Buneman P., Burgess N., Canals S., Carloni P., Carsten Stahl B., Chandran S., Changeux J.-P., Clasca F., Coenen C., Costa R., Curioni A., D'Angelo E., Dal Peraro M., Dartigues J.-F., Davidson S., Davison A., Deco G., De Estudios A., DeFelipe J., Dehaene S., Dehaene G., De Los Rios P., De Schutter E., Destexhe A., De Zeeuw C., Diesmann M., Draganski B., Dudai Y., Dürr A., Ehrmann O., Eicker N., Eickhoff S., Einevoll G., Erbacci G., Ertl T., Evans A., Evers K., Faugeras O., Fischer H., Frackowiak R., Frégnac T., Freund T., Fries P., Frisoni G., Frommer A., Fua P., Furber S., Gamrat C., Gehrke J., Gerstner W., Gewaltig M.-O., Giaume C., Giese M., Girona S., Goebel G., Gottfried F., Grant K., Grant S., Griebel M., Grillner S., Grimes K., Grollier J., Grübl A., Grubmüller H., Grün S., Hagoort P., Halevy M., Hamaekers J., Hardt M., Hari R., Hausser M., Hautop Lund H., Hellgren-Kotaleski J., Herzog M., Hess Bellwald K., Hill S., Hille K., Hines M., Huret A., Husmann D., Ioannidis Y., Jérusalem A., Jirsa V., Jonas P., Karni A., Kersten M., Keyes D., Kherif F., Kindler B., Kleinfeld D., Klinker G., Klöppel S., Klüver L., Knoll A., Koch Christof, Koch Christoph, Kopanitsa M., Kuhlen T., Labarta J., Laio A., Lansner A., Laure E., Laurent G., Lavery R., Le Bihan D., Leblebici Y., Lester D., Lindahl E., Lippert T., Lujan R., Luthi-Carter R., Maass W., Macii E., Maestu F., Magistretti P., Mainen Z., Malach R., Mangin J.-F., Marcus-Kalish M., Marin O., Markram H., Marris C., Martinez S., Martone M., Mayr C., Meier K., Menasalvas E., Michelle C., Miglino O., Migliore M., Mohammed A., Mohr B., Morrison A., Nieder A., Niederberger R., Noé F., Nolfi S., Nyberg L., Orgogozo J.-M., Orozco M., Orth B., Owen M., Owen R., Ozguz V., Pallier C., Parkkonen L., Pastor L., Pavone F., Paz R., Peña JM., Pessiglione M., Petkov N., Pinel P., Pocklington A., Poline J.-B., Ponting C., Pottmann H., Poupon C., Ramirez A., Reitmayr G., Renaud S., Robbins T., Robles V., Ros E., Rose N., Roth A., Röthlisberger U., Roweth D., Rückert U., Saksida L., Salles A., Sanchez-Vives M., Sandi C., Sansom M., Sanz R., Saria A., Scattoni ML., Schemmel J., Schneider F., Schrader S., Schrauwen B., Schüffny R., Schulthess T., Schürmann F., Segev I., Senn W., Shanahan M., Shillcock J., Sigman M., Singer W., Slater M., Smit A., Smith LS., Soriano E., Spijker S., Spiliopoulou M., Sporns O., Stampanoni A., Steinmacher-Burow B., Suarez E., Taly A., Tarek M., Tass P., Thill S., Thirion B., Thompson P., Thomson A., Tiesinga P., Toga A., Trappl R., Triller A., Tsodyks M., Ullmann S., Valero M., van der Smagt P., van Leeuwen C., van Wassenhove V., Vidal F., Villani C., Vincente M., Voet T., von Landesberger T., Wade R., Wang Y., Weber B., Weiskopf N., Wierstra D., Williams R., Wittum G., Wolf F., Wolkersdorfer K., Yokoyama C., Ziemke T., Zilken H.

Graphics design: Marie-Eve Laurent

Table of contents

Abstract	3
EXECUTIVE SUMMARY	7
The Human Brain Project	8
A new foundation for brain research	9
Society and ethics	12
Leveraging the strengths and diversity of European research	12
Three phases	12
Cost	13
Governance and management	13
Impact	13
REPORT	15
1	
Why this report?	16
2	
Vision and rationale	17
Why the brain?	17
Why now?	18
A new foundation for brain research	20
Data	20
Theory	22
ICT platforms	22
Applications	25
Why a FET Flagship?	27
3	
Science and technology plan	28
Data	28
Theory	35
ICT platforms	37
Applications	48
Coordination of resources and research communities	55
4	
Implementation	57
The challenge	57
Research organisation	57

Governance	58
Management	61
Costs and revenues	61
Additional financial resources	64
Know-how and technical facilities	64
Initial HBP Consortium	66
Gender equality	73
Alignment with national, European and international priorities	73
Use of results and dissemination of knowledge	79
Risks and contingencies	82
5	
Impact	84
Science	84
Medicine	86
Future computing technology	88
The European economy and industry	90
Society and ethics	91
CONCLUSIONS	95
Glossary	98
References	101

Executive summary

The Human Brain Project

Understanding the human brain is one of the greatest challenges facing 21st century science. If we can rise to the challenge, we can gain fundamental insights into what it means to be human, develop new treatments for brain diseases, and build revolutionary new Information and Communications Technologies (ICT). In this report, we argue that the convergence between ICT and biology has reached a point at which it can turn this dream into reality. It was this realisation that motivated the authors to launch the Human Brain Project – Preparatory Study (HBP-PS) – a one year EU-funded Coordinating Action in which nearly three hundred experts in neuroscience, medicine and computing came together to develop a new “ICT-accelerated” vision for brain research and its applications. Here, we present the conclusions of our work.

We find that the major obstacle that hinders our understanding of the brain is the fragmentation of brain research and the data it produces. Modern neuroscience has been enormously productive but unsystematic. The data it produces describes different levels of biological organisation, in different areas of the brain in different species, at different stages of development. Today we urgently need to integrate this data – to show how the parts fit together in a single multi-level system.

Thanks to the convergence between biology and ICT, this goal is within our grasp. New sequencing and imaging technologies and new techniques of microscopy have revolutionised our ability to observe the brain. Cloud technology, combined with the Internet, allows us to federate data from research groups and clinics all over the world, Neuroinformatics gives us new means to analyse the data, to build and share detailed brain atlases, to identify gaps in our knowledge and to predict the value of parameters where experimental data is still missing. Supercomputers make it possible to build and simulate brain models with unprecedented levels of biological detail.

These technologies can enormously accelerate brain research. They can also open the road to treatments that prevent and cure brain disease and to new computing technologies with the potential to revolutionise industry, the economy and society. Medical informatics can mine enormous volumes of clinical data allowing us to understand the basic causes of brain diseases, a pre-condition for early




Figure 1: The Human Brain - one of the greatest challenges for 21st century science

diagnosis, prevention and cure. Meanwhile, brain simulation has the potential to revolutionise ICT itself, laying the foundations for a completely new category of low-energy computing systems, with brain-like intelligence. If European industry is to play a leading role in the world economy of the 2020s and 2030s, it has to take the lead in developing these technologies.

Applying ICT to brain research and its applications promises huge economic and social benefits. But to realise these benefits, we first have to make the technology accessible to scientists, building it into research platforms they can use for basic and clinical research, drug discovery, and




Figure 2: HBP research areas

technology development. In brief, ICT-based brain science has to become a strategic area for European research.

To set this effort in motion, we need a catalyst. In this report, we propose that this should take the form of an integrated system of ICT-based research platforms providing radically new services to neuroscientists, clinical researchers and technology developers. Such platforms, we will argue, would create a new foundation for brain research, enormously accelerating the translation of basic science results into concrete benefits for European health services, the health of European citizens and the competitive position of European industry.

Building and operating the platforms is technically feasible. The necessary supercomputing technology is developing rapidly, in step with a well-established industry road-map. Prototypes of other key technologies have already been tested in the labs of potential partners. What is needed is a clear vision, strong, flexible leadership, long-term investment in research and engineering and a strategy that leverages the diversity and strength of European research. These requirements can only be met by a project on the scale of a FET Flagship. In this report, therefore, we propose that the European Commission launches such a Flagship. We call it *The Human Brain Project* (HBP). We summarise the goal of the project as follows.

The Human Brain Project should lay the technical foundations for a new model of ICT-based brain research, driving integration between data and knowledge from different disciplines, and catalysing a community effort to achieve a new understanding of the brain, new treatments for brain disease and new brain-like computing technologies.

A new foundation for brain research

The Human Brain Project should pursue four goals, each building on existing work, and acting as a catalyst for new research.

1. *Data*: generate strategically selected data essential to seed brain atlases, build brain models and catalyse contributions from other groups.
2. *Theory*: identify mathematical principles underlying the relationships between different levels of brain organisation and their role in the brain's ability to acquire, represent and store information.
3. *ICT platforms*: provide an integrated system of ICT platforms offering services to neuroscientists, clinical researchers and technology developers that accelerate the pace of their research.
4. *Applications*: develop first draft models and prototype technologies, demonstrating how the platforms can be used to produce results with immediate value for basic neuroscience, medicine and computing technology.

Data

Modern neuroscience research has already generated huge volumes of experimental data; large-scale initiatives already in progress will produce a deluge of new findings. Even then, however, much of the knowledge needed to build multi-level atlases and unifying models of the brain will still be missing. The first goal for the HBP should thus be to generate and interpret strategically selected data, unlikely to come from other sources. The HBP-PS has identified three main focuses for this research.

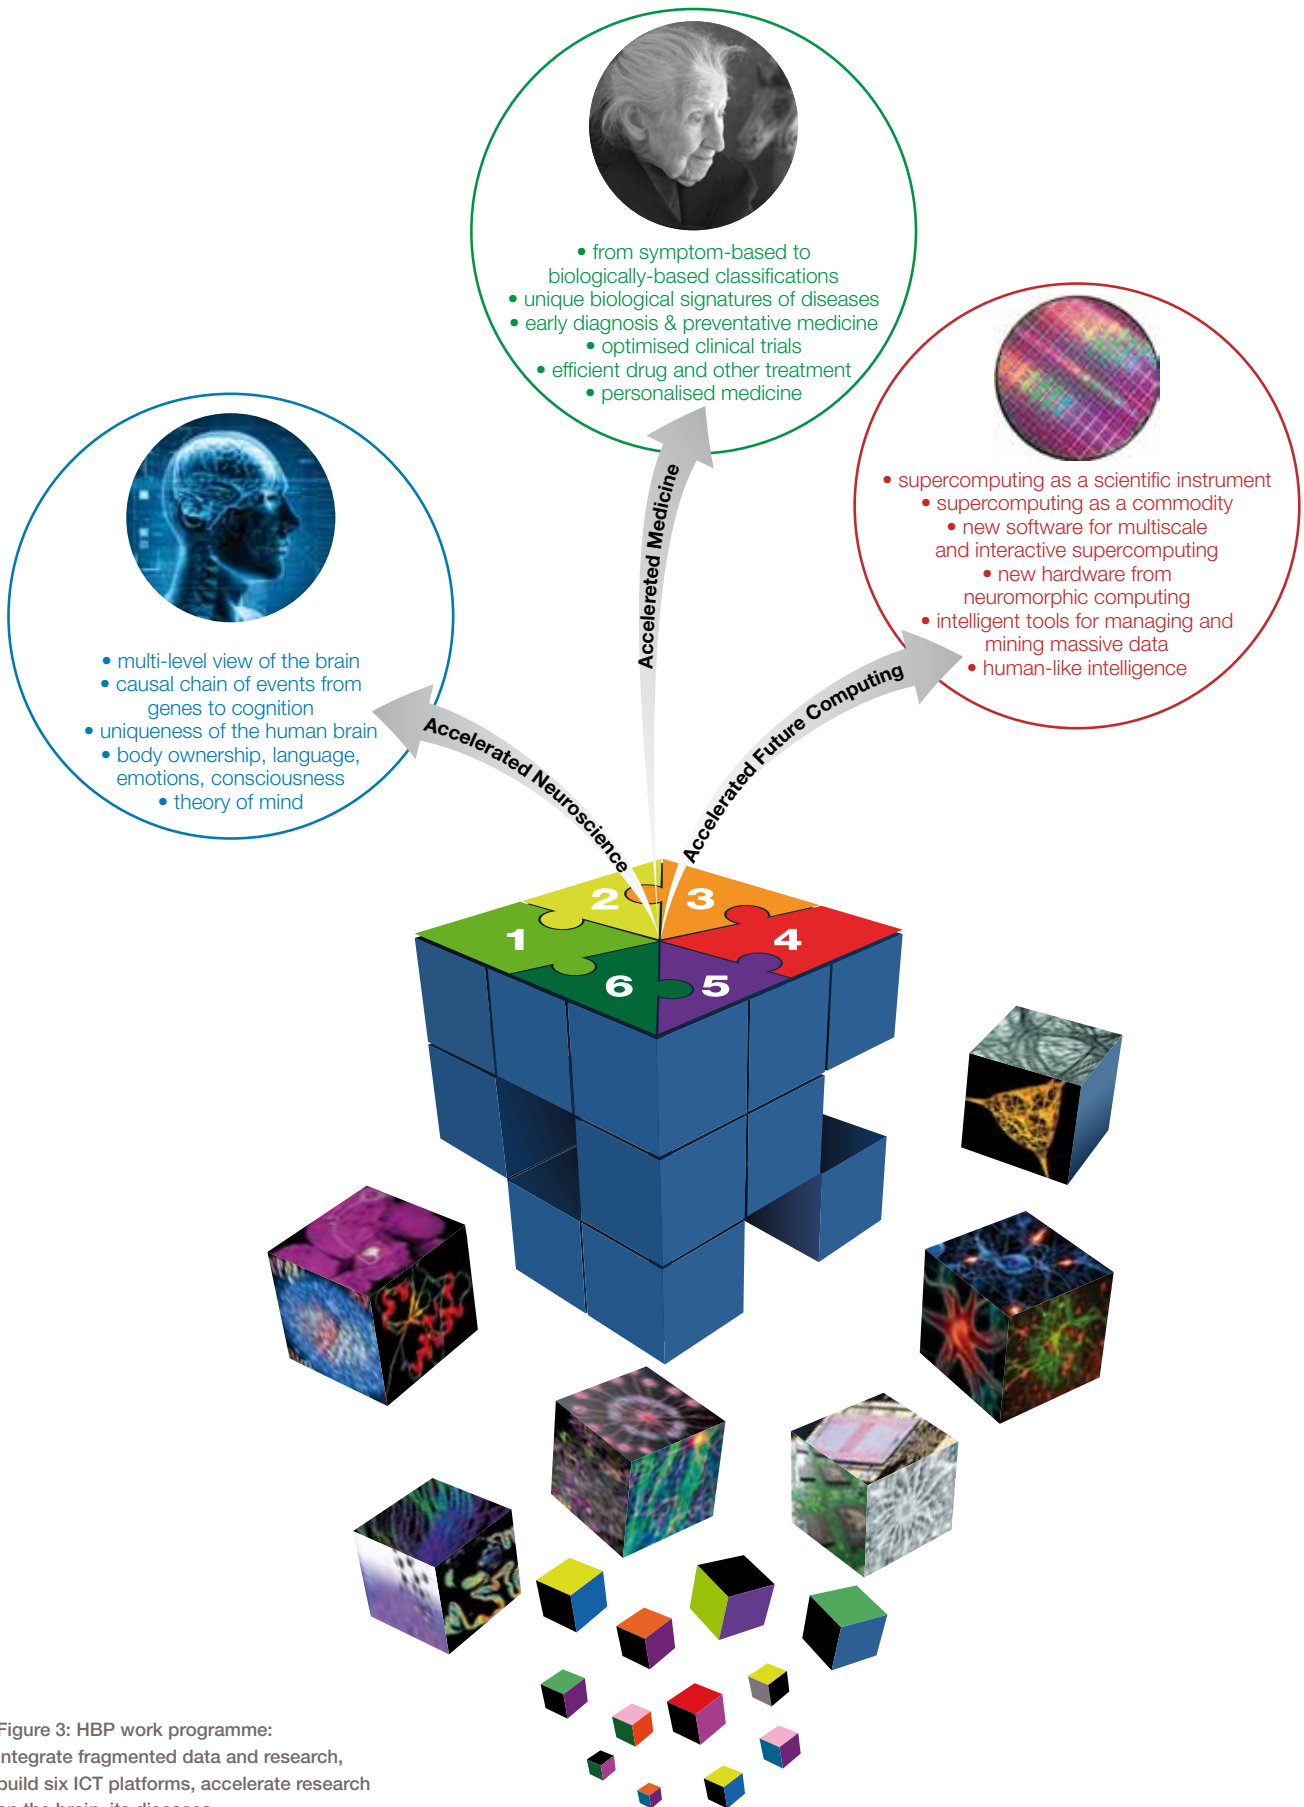


Figure 3: HBP work programme: integrate fragmented data and research, build six ICT platforms, accelerate research on the brain, its diseases and future computing technologies

Multi-level brain structure in mouse. Many results from studies of the mouse brain are applicable to all mammals. A systematic study of the relations among its different levels of organisation would provide vital input for atlases and models of the human brain.

Multi-level structure of the human brain. Mouse data provides many insights into the human brain. Obviously, however, the human brain is different. To identify and characterise these differences, HBP research should generate strategically selected data for the human brain, as far as possible matching the data available for mouse.

Brain function and neuronal architectures. One of the HBP's most important goals must be to understand the relationship between brain structure and function. A third focus for HBP research should thus be the neuronal architectures responsible for specific cognitive and behavioural skills – from simple capabilities, also present in non-human species, to those such as language, that are exclusive to humans.

Theory

Without sound theoretical foundations, it will not be possible to overcome the fragmentation of neuroscience data and research. The HBP should thus include a concerted programme of theoretical research, focusing on the mathematical principles underlying the relationships between different levels of brain organisation and the way the brain acquires, represents and stores information. As part of this programme, the HBP should establish a *European Institute for Theoretical Neuroscience*, encouraging participation by scientists from outside the project and acting as an incubator for novel approaches.

ICT platforms

The HBP's third goal should be to create an integrated system of ICT platforms with the potential to set in motion a new kind of ICT-based brain research. We propose that there should be six of these platforms, dedicated to Neuroinformatics, Brain Simulation, Medical Informatics, High Performance Computing, Neuromorphic Computing, and Neurorobotics.

Neuroinformatics. The HBP *Neuroinformatics Platform* should provide technical capabilities making it easier for neuroscientists to analyse structural and functional brain data and to build and navigate multi-level brain atlases. The platform should also include tools for *predictive neuroinformatics*, making it possible to detect statistical regularities in the relationships between data representing different levels of brain organisation and to estimate the values of parameters that are difficult or impossible to measure experimentally. These tools will provide a new way of filling the gaps in data and knowledge that current-

ly prevent us from achieving an integrated understanding of the brain.

Brain Simulation. The HBP should create a large-scale *Brain Simulation Platform*, making it possible to build and simulate multi-scale brain models at the different levels of detail appropriate to different scientific questions. The platform – which would play a central role in the whole project – should provide researchers with modelling tools, workflows and simulators allowing them to integrate large volumes of heterogeneous data in multi-scale models of the mouse and human brains, and to simulate their dynamics. This possibility would allow them to perform *in silico* experiments, impossible in the lab. Tools provided by the platform would generate input essential for HBP research in medicine (models of diseases and the effects of drugs), neuromorphic computing (brain models for implementation in neuromorphic hardware), and neurorobotics (models of neural circuitry for specific cognitive and behavioural tasks).

High Performance Computing. The HBP *High Performance Computing Platform* should provide the project and the community with the computing power they need to build and simulate models of the brain. This should include both the latest, most powerful supercomputing technology, up to the exascale, and completely new capabilities for interactive computing and visualisation.

Medical Informatics. The HBP *Medical Informatics Platform* should federate clinical data from hospital archives and proprietary databases, while providing strong protection for patient data. Such capabilities would allow researchers to identify “biological signatures” for specific disease processes – a fundamental breakthrough. Once researchers have an objective, biologically-grounded way of detecting and classifying diseases, they will be able to understand their causes and develop effective treatments.

Neuromorphic Computing. The HBP *Neuromorphic Computing Platform* should provide researchers and application developers with the hardware and design tools they need to develop systems and devices modelled on the architecture of the brain and prototype applications. The new platform would allow researchers to develop a completely new category of compact, low-power devices and systems approaching brain-like intelligence.

Neurorobotics. The HBP *Neurorobotics Platform* should provide researchers with tools and workflows allowing them to interface detailed brain models to a simulated body in a simulated environment, comparing the behaviour they can learn against results from human and animal experiments. These capabilities would offer neuroscientists new strategies for studying the multi-level mechanisms underlying behaviour. From a technological perspective it would give developers the tools they need to develop robots with the potential to achieve human-like capabilities, impossible to realise in systems that do not have a brain-like controller.

Applications

The project's fourth major goal should be to demonstrate the value of its platforms for fundamental neuroscience research, for clinical studies and for technology development. We expect that successful demonstrations would trigger a wave of research by groups outside the project. To encourage this effect, a large proportion of this work should be entrusted to groups not included in the initial HBP Consortium. Later in this report, we will outline specific proposals to achieve this goal.

Integrative principles of cognition. Researchers should use the *Brain Simulation* and *Neurorobotics Platforms* in projects that systematically dissect the neuronal circuits responsible for specific behaviours, simulating the effects of genetic defects, lesions, and loss of cells at different levels of brain organisation and modelling the effects of drugs. The ultimate goal should be to model the unique capabilities that distinguish humans from other animals, in particular language. Such models would represent a fundamental advance in our understanding and would have immediate applications in medicine and technology.

Understanding, diagnosing and treating brain disease. Research should exploit the capabilities of the *Medical Informatics*, *Neuroinformatics* and *Brain Simulation Platforms* to discover biological signatures associated with specific disease processes, to understand and simulate these processes, and to identify new targets for prevention and treatment. This work should demonstrate the ability of the HBP platforms to produce immediately valuable results. New diagnostic tools would make it possible to diagnose disease earlier, before it causes irreversible damage, develop new drugs and test new treatment strategies adapted to the needs of specific patients – so-called personalised medicine. The end result would be better outcomes for patients and lower healthcare costs. Better understanding and diagnosis would also help to optimise the drug discovery process, allowing better screening of drug candidates and better selection of patients for clinical trials. The benefits would include a reduction in expensive failures during late trials and reductions in the cost of developing new drugs, currently estimated at around Eur 1 billion per drug.

Future Computing Technologies. Researchers should use the HBP's *High Performance Computing*, *Neuromorphic Computing* and *Neurorobotics Platforms* to develop new computing technologies and new applications. The *High Performance Computing Platform* would allow them to design hybrid technology integrating neuromorphic devices with conventional supercomputing. With the *Neuromorphic Computing* and *Neurorobotics Platforms*, they would be able to build prototypes of applications with large potential markets. These would include robots for use in the home, manufacturing and services, as well as “invisible”, yet equally significant technologies, such as data mining and controllers for vehicles, household appliances, manufacturing, image and video processing, and telecommunications.

Society and ethics

Given the large potential impact of HBP research and technology, it is essential that the project follow a policy of Responsible Innovation. The HBP should thus include a far-reaching Society and Ethics Programme, funding academic research into the potential social and economic impact of HBP research, and its ethical and conceptual implications, managing programmes to raise ethical and social awareness among HBP researchers, and, above all, encouraging an intense dialogue with stakeholders and with civil society that gives full expression to inevitable differences in approaches and values.

Leveraging the strengths and diversity of European research

The HBP should leverage the strengths and diversity of European research by including scientists from different schools of thought with a broad range of theoretical and experimental approaches. We therefore recommend that the project dedicate a steadily rising proportion of its annual funding to research by scientists from outside the initial HBP Consortium, with calls for proposals beginning during the ramp-up phase. By the end of the project at least 40% of the annual budget should go to outside researchers. We recommend that this funding should be provided through two new programmes. The first should focus on individual researchers, following the example of the ERC and Marie-Curie grant schemes. The second, modelled on the FET programme, should focus on research groups proposing their own projects to use the HBP platforms.

Three phases

We propose that the HBP should be organised in three phases, lasting a total of ten years.

For the first two and a half years (the “ramp-up” phase), the project should focus on setting up the initial versions of the ICT platforms and on seeding them with strategically selected data. At the end of this phase, the platforms should be ready for use by researchers inside and outside the project.

For the following four and a half years (the “operational phase”), the project should intensify work to generate strategic data and to add new capabilities to the platforms, while simultaneously demonstrating the value of the platforms for basic neuroscience research and for applications in medicine and future computing technology.

In the last three years (the “sustainability phase”), the project should continue these activities while simultaneously moving towards financial self-sustainability – ensuring that the capabilities and knowledge it has created become a permanent asset for European science and industry.

Cost

We estimate that the total cost of the HBP would be approximately **Eur 1,190 million**, of which **Eur 80 million** for the ramp-up phase, **Eur 673 million** for the operational phase and **Eur 437 million** for the sustainability phase. The funding required from the EU Commission would amount to approximately **Eur 643 million**.

Governance and management

The Human Brain Project would be a large, ten-year interdisciplinary project, involving partners from more than twenty countries and a large budget. It is essential, therefore, that the project's governance and management provide strong, flexible leadership, while simultaneously guaranteeing that the ICT platforms become a genuine community resource. *Mission critical activities* should be carried out by a consortium of partners (the *HBP Consortium*) whose composition would evolve over the duration of the project. *Research using the ICT platforms* should take the form of research projects, selected through a competitive process open to the entire scientific community.

Impact

The Human Brain Project would enormously accelerate progress towards a multi-level understanding of brain structure and function, towards better diagnosis, better understanding, and better treatment of brain diseases and towards new brain-inspired Information and Communications Technologies. The impact on European science, European industry, the European economy and European society is potentially very large.

Scientifically speaking, the data generated by the HBP and the technological capabilities offered by the project's ICT platforms would help to overcome the fragmentation of neuroscience research, opening the door to a completely new understanding of the relationships between brain structure and function. The project would allow researchers to address some of the most important challenges facing modern neuroscience, including learning and memory, the nature of the so-called neural code, and even the neuronal mechanisms of consciousness and awareness.

The Human Brain Project would also make a huge impact in medicine, accelerating the development of better diagnostic tools, and better treatments. Given the huge cost of brain disease, even small improvements (e.g. earlier diagnosis, therapies that delay cognitive decline in neurodegenerative disease, etc.) would produce large economic and social benefits. Reductions in the cost of drug discovery, and improvements in success rates, would provide important benefits for the pharmaceutical industry. The key modelling

and simulation know-how would be developed by European researchers and institutions, strengthening the competitive position of the industry in the development of new drugs for brain disease – a potentially enormous market.

By integrating brain research with ICT, the HBP would help to determine the future direction of computing technology. In supercomputing, new techniques of interaction and visualisation, multi-scale simulation and cloud computing developed by the project, would stimulate the development of new services for industry and consumers, initiating a virtuous circle in which increased demand leads to economies of scale and falling costs, and falling costs further boost demand, making supercomputers available to far wider sectors of academia and industry than at present. Many of these capabilities would depend on advanced software – an area in which Europe has a strong competitive advantage.

HBP work in neuromorphic computing and neurobotics would open the road for the development of compact, low-power systems with the long-term potential to achieve brain-like intelligence. Although these technologies will not replace the “conventional” computing technologies that have driven the last fifty years of European growth, the range of their potential applications and their strategic importance are on the same scale. By taking the lead in their development, the HBP would play a vital role in securing Europe's competitive position in the world economy.

“The appeal of neuromorphic architectures lies in i) their potential to achieve (human-like) intelligence based on unreliable devices typically found in neuronal tissue, ii) their strategies to deal with anomalies, emphasizing not only tolerance to noise and faults, but also the active exploitation of noise to increase the effectiveness of operations, and iii) their potential for low-power operation. Traditional von Neumann machines are less suitable with regard to item i), since for this type of tasks they require a machine complexity (the number of gates and computational power), that tends to increase exponentially with the complexity of the environment (the size of the input). Neuromorphic systems, on the other hand, exhibit a more gradual increase of their machine complexity with respect to the environmental complexity. Therefore, at the level of human-like computing tasks, neuromorphic machines have the potential to be superior to von Neumann machines.”

(source: International Technology Roadmap for Semiconductors, 2011)

Report

1

Why this report?

This report summarises the results of the *Human Brain Project – Preparatory Study*, an EU Coordinating Action in which the authors – nearly three hundred experts in neuroscience, medicine and computing – worked together to develop a new vision for brain research and its applications. We conclude that turning this vision into reality will require a project on the scale of a FET Flagship. We therefore propose that the European Commission launches such a Flagship, which we call *The Human Brain Project*. We summarise the goal of the project as follows.

The Human Brain Project should lay the technical foundations for a new model of ICT-based brain research, driving integration between data and knowledge from different disciplines, and catalysing a community effort to achieve a new understanding of the brain, new treatments for brain disease and new brain-like computing technologies.

The remainder of this report sets out the grounds for our recommendation. The second part (“Vision and rationale”) puts the argument why brain research should be a priority for ICT, why we should invest now rather than later and why achieving our goals requires an initiative on the scale of a FET Flagship. In the third part (“Science and technology plan”), we analyse issues of scientific and technical feasibility, describing precise objectives and the underlying science and technology, identifying critical problems and feasible solutions, and explaining what we can expect to achieve in terms of data, theory, ICT platforms and applications. In part four, we discuss implementation: the way a FET Human Brain Project should be governed, the availability of know-how and technical resources, the estimated cost and the possible risks. Finally, in part five, we examine the impact of the project and its potential benefits for European science, European industry and European citizens.

2 Vision and rationale

Why the brain?

The human brain participates in every human emotion, every human feeling, every human thought and every human decision. No other natural or engineered system can match its ability to adapt to novel challenges, to acquire new information and skills, to take complex decisions and to work reliably for decades on end. And despite its many diseases, no other system can match its robustness in the face of severe damage or match its amazing energy efficiency. Our brain consumes about 30W, the same as an electric light bulb, thousands of times less than a small supercomputer.

The human brain is a massively complex information processing system with a hierarchy of different yet tightly integrated levels of organisation: from genes, proteins, synapses and cells to microcircuits, brain regions, and the whole brain (see Figure 4). Today, we know a lot about the individual levels. What we do not have is a causal understanding of the way events at the lowest level in the hierarchy cascade through the different levels to produce human cognition and behaviour. For example, more than a hundred years of research has yet to give us a proper understanding of the link from synaptic plasticity to learning and memory, or of the way a gene defect works through the different levels to produce disease. Achieving this kind of understanding is a major challenge for neuroscience with implications that go far beyond research: if we could understand the brain we could prevent or cure brain diseases such as autism, depression and Alzheimer's; we could also produce new computing technologies that share the brain's ability to operate reliably on very little power, and its ability to learn.

Medical research has identified over five hundred brain diseases, ranging from migraine and addiction to depression and Alzheimer's. An authoritative study has estimated that in 2010, more than a third of European citizens were directly affected by at least one of these diseases. The same study estimated the cost to the European economy at nearly Eur 800 billion [1]. As European populations age, the number of citizens affected and the cost of their care will inevitably grow, potentially to unsustainable levels.

Today, these diseases are usually diagnosed in terms of symptoms and syndromes, an approach that makes it very difficult to produce correct diagnoses, or even to select patients for clinical trials. To prevent and cure brain disease,

Spatial scales

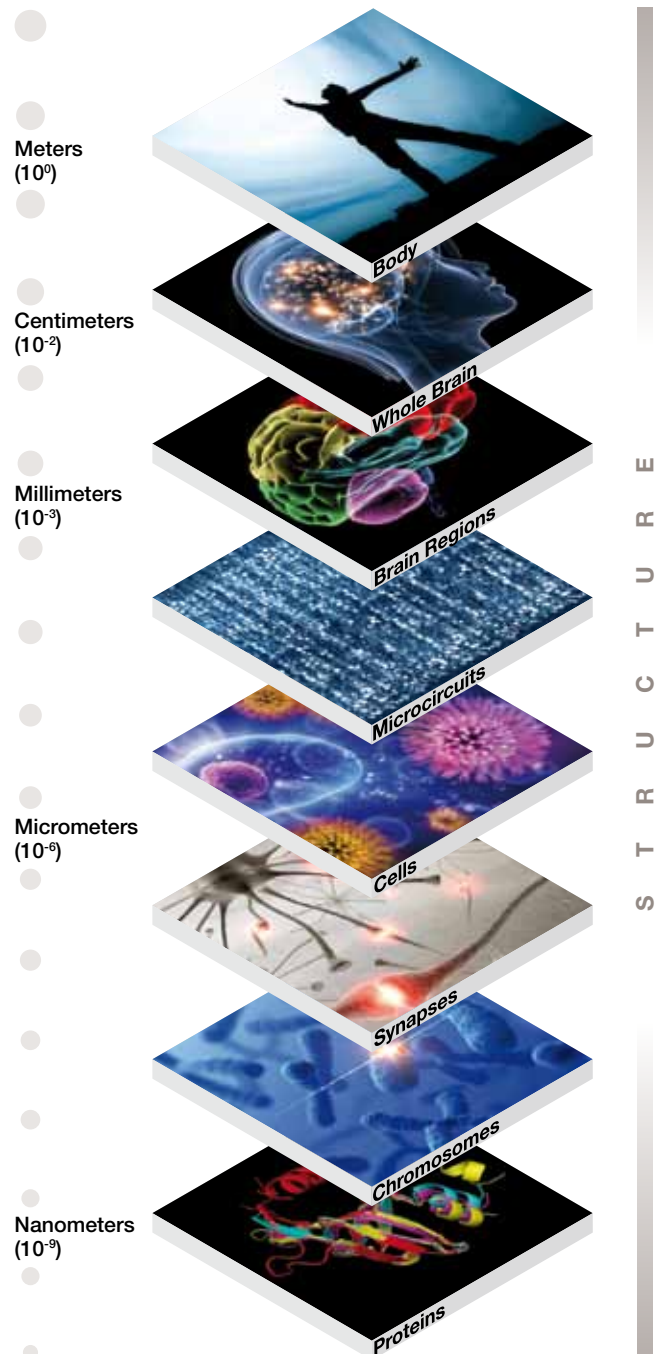


Figure 4: From molecules to the body: spatial scales for the brain's different levels of organisation span nine orders of magnitude

Neuroscience, medicine, and ICT all require an integrated multi-level understanding of the brain

researchers need to understand their underlying causes. Studies have shown that conditions classified as a single disease can have completely different genetic causes and that diseases with different causes can produce very similar symptoms. This means we need to move beyond isolated studies of individual disorders and investigate brain diseases systematically, identifying key similarities and distinctions, creating new biologically grounded classifications, and understanding the complex disease processes leading from changes at lower levels of brain organisation to clinical symptoms. Only then will it be possible to intervene effectively. Today, the knowledge we need is lacking.

Less obviously, poor understanding of the brain is also an obstacle for ICT. The computing technologies that have driven world economic growth since the 1950s are rapidly approaching fundamental limits on processing speed, power consumption, reliability and programmability. This creates an urgent need for new architectures – perhaps also for completely new computing paradigms. According to the experts responsible for the International Technology Roadmap for Semiconductors [2], the most effective strategy for moving forward is to complement current information technology with compact, low-power computing systems inspired by the architecture of the brain, with human-like intelligence. But to actually implement such systems, we first have to reach a deep understanding of the brain's cognitive architectures: the way in which the different levels of organisation work together to acquire, represent, and store information.

In summary, neuroscience, medicine, and ICT all require an integrated multi-level understanding of the brain.

Why now?

An integrated multi-level understanding of the brain would offer enormous benefits for medicine and for future computing technology – indeed for the European economy and European society as a whole. It seems obvious that if it is possible to achieve such an understanding, it should be a top priority for research. In this report, we will argue that today, for the first time, technology has reached a point when it is possible.

The last twenty years have seen a rapid convergence between Information and Communications Technologies and the life sciences that has revolutionised our ability to observe and understand the brain (see Figure 5). Exponential improvements in sequencing technology have made it feasible

to sequence whole genomes, rapidly, at ever-lower costs. New techniques of ultra-microscopy can create detailed 3D pictures of whole animal brains, tracing the circuits involved in specific functions. Rapid improvements in imaging technology have made it possible to image the structure of the brain at ever higher resolution, to map the fibres linking different areas, and to elucidate the neuronal circuitry responsible for specific functions. Cloud technology, combined with the Internet and modern cryptography, allows us to federate data from research groups and clinics all over the world, while simultaneously providing strong protection for patient data. Medical informatics allows us to “mine” this data for biological signatures of disease, providing clues to the causes of disease, better diagnosis, personalised treatment and new targets for drug discovery. Neuroinformatics gives us the tools to build detailed brain atlases, and to share them with the community. Other neuroinformatics tools can mine massive volumes of data for recurrent patterns, predict missing values and fill in gaps in our knowledge. High performance computing has given us the computing power to build and simulate multi-scale brain models with unprecedented levels of biological detail. Simulation technology allows us to connect brain models to virtual agents interacting with virtual environments, enabling measurements and manipulations impossible in the biology lab. In computing, sub-micron technologies and many-core technologies allow us to implement brain models in compact, low-power computing devices with a massive range of potential applications, from robots to industrial machinery.

Taken together, these technologies can provide a new ICT-based foundation for brain research, medicine and future computing technology.

European groups are already world leaders in many of the relevant areas of science and technology. For instance, the Wellcome Trust – Sanger Institute played a vital role in the Human Genome Project and in the development of modern genomics and sequencing technology. The French NeuroSpin centre is playing a leading role in the development and applications of neuroimaging. Europe plays a leading role in the International Neuroinformatics Coordinating Facility (INCF), which has its headquarters in Stockholm. The Swiss Blue Brain Project has pioneered detailed biological simulation of the brain. European consortia such as PRACE are driving the development of exascale computing. European-funded research projects such as FACETS, Brain-ScaleS and ECHORD and nationally funded initiatives such as SpiNNaker are leading the development of neuromorphic computing systems.

Today, a European-led Human Brain Project in the FET Flagship Programme would reinforce this leadership and help the transfer of basic research knowledge to industry. However, there is a limited window of opportunity. Neuroscience and ICT are progressing rapidly. If Europe does not take the lead, others will. In the world economy of the 2020s and the 2030s, pharmacological treatment for brain disease and brain-inspired computing will be key drivers of economic growth. If Europe wants to establish a strong competitive position in these sectors, the time to act is now.

ICT

NEUROSCIENCE

MEDICINE

COMPUTING

Allen Institute launches \$300 million mouse brain project and commits \$1B for brain research	2012	Whole exome sequencing and informatics reveals new genetic key to autism	2012	More than 612 million websites; more than 900 million Facebook users
One Mind for Research; \$6B targetted funding; Active transposons in the human brain	2011	> 350 biotechnology-based products resulting from Human Genome Project in clinical trial	2011	IBM Watson wins Jeopardy; BrainScaleS; ARM shipments exceed 30 billion; IBM Neuro Chip
NIH Human Connectome Project; MIT Intelligence Initiative	2010	Cancer Genome Project; Synthetic cell	2010	PRACE, Brain-I-Nets, DARPA Neovision II
First single cell transcriptome	2009	Large genome-wide association studies identify new Alzheimer's disease genes	2009	HPC driven real-time realistic global illumination algorithm
Janelia Farms founded; NIH Open Access Policy to enhance knowledge dissemination	2008	First Eukarya interactome; deep brain stimulation to move muscles	2008	Cray XT5 Petaflop on superconductive materials; DARPA SYNAPSE project started
Decade of the Mind	2007	Visual prosthetics, HT-DNA sequencing; embryonic stem cells from human skin	2007	Apple iPhone
Optogenetics technology	2006	First pluripotent stem cells; 23andMe launch personal genomics; BrainGate System for paraplegics	2006	Amazon Cloud, SpiNNaker Project; real-time animation of digital human bodies
International Neuroinformatics Coordinating Facility (INCF) founded; Blue Brain Project launched	2005	HapMap & First Genome Wide Association Study; identification of human microRNA genes	2005	FACETS, DAISY, COLAMN, DARPA Aug Cognition; IBM Cell Processor; Assisted GPS for cell phones
Champalimaud Foundation launches new brain research institute, commits €500M	2004	ADNI founded; comparative genomic analysis identifies cause for Bardet-Biedl syndrome	2004	INTEL Dual Core; DARPA Neovision I; Facebook triggers social networking phenomenon
Paul Allen founds the Allen Institute to map the mouse transcriptome; commits \$100M	2003	Human genome completed	2003	About 1 billion PCs sold
	2002		2002	Earth simulator
NEST Large-scale neural network simulator	2001	Telesurgery; 1000 Genomes Project; artificial liver	2001	
Institute of Systems Biology	2000	First draft human and mouse genome; human testing of an Alzheimer's disease vaccine	2000	
	1999	Structural Genomics Project; fruit fly sequenced	1999	
Crystal structure of an ion channel; Gatsby Comp. Neurosc. Institute launched. Lord Sainsbury commits £500M	1998	Stem cell therapy	1998	Cray T3E Teraflop modelling metallic magnets; MDA Silicon Brain Program
	1997	Protein accumulation in human neurodegenerative diseases	1997	Google founded
	1996	Dolly, the sheep cloned	1996	MDA Silicon Neuron Program
	1995	First bacterial genome	1995	Real-time image-based rendering; JAVA; Support Vector Machines
	1994	First HIV (anti-AIDS) treatment; gene therapy via implanted transformed fibroblasts	1994	WWW Foundation launched; GPS goes live with 24 satellites
Int. Consortium for Brain Mapping; NIH Human Brain Project	1993	Huntington's disease gene identified; interferon for multiple sclerosis	1993	50 websites in the world
First whole head MEG system; fMRI	1992	Self-folding prion proteins	1992	
	1991	Amyloid hypothesis of Alzheimer's Disease; term "evidence-based medicine" coined	1991	Motion capture technology for guiding virtual robots
Two-photon microscope US "Decade of the brain"	1990	First human viral gene therapy	1990	WWW technology
DNA microarray technology for transcriptomics	1989		1989	INTEL Electrically Trainable Artificial Neural Network chip
	1988		1988	Silicon retina published; 45 million PCs in the USA; TAT-8, first transatlantic fiber optic cable
	1987	Deep-brain electrical stimulation; laser surgery on human cornea; meningitis vaccine developed	1987	Cray YMP GigaFlop on finite element analysis Thinking Machine with 65536 processors
Mapping of the structure of C.elegans nervous system	1986		1986	Radiosity rendering for scene illumination
	1985	TMS -Transmagnetic stimulation; automated DNA sequencer; surgical robots	1985	First GigaFlop supercomputer (uses ARM processors)
Transgenic mouse produced; Boltzman machine	1984		1984	
Cylindric PET scan; Polymerase Chain Reaction technology (PCR)	1983	Automated DNA sequencers (ABI 380A); discovery of HIV	1983	
The Hopfield Artificial Neural Network; Statistical Mechanics of Artificial Neural Networks	1982	Interferon cloning; commercial protein sequencers	1982	First 32 bit microprocessor
Magnetic Resonance Imaging (MRI)	1981	Artificial Skin; Applied Biosystems founded by two HP engineers	1981	IBM Personal Computer
	1980	WHO declares smallpox eradicated	1980	Recursive ray tracing and new algorithms for realistic scene rendering
	1979	Anti-viral drugs	1979	
Patch Clamp technique	1978	Test-tube baby born; first cochlear implant	1978	
First Molecular Dynamics simulation of a protein	1977	First DNA sequencing method; first virus sequenced	1977	

Figure 5: The merging of neuroscience, medicine and computing. Key events 1977-2012

A new foundation for brain research

We recommend that a Human Brain Project should pursue four goals, each building on existing work, and acting as a catalyst for new research.

1. *Data*: generate strategically selected data (data on the structure of the mouse and human brains, data on human brain function) essential to catalyse integration among existing data sets, seed brain atlases and build brain models.
2. *Theory*: identify mathematical principles underlying the relationships between different levels of brain organisation and their role in acquiring, representing and storing information about the outside world.
3. *ICT platforms*: provide an integrated system of ICT platforms, allowing researchers to federate and analyse massive volumes of heterogeneous neuroscience and clinical data, build and simulate multi-scale models of the brain, couple these models to simulated and physical robots, and use them as the basis for radically new computing technology.
4. *Applications*: fund research projects that use the platforms to accelerate basic neuroscience (dissecting the biological mechanisms responsible for cognition and behaviour), medicine (understanding brain disease, finding new treatments, personalised medicine) and technology development (low-energy computing systems with brain-like intelligence, hybrid systems integrating neuromorphic and conventional technologies, new applications for industry, services, vehicles, the home).

Data

Modern neuroscience research has generated vast volumes of experimental data and large-scale initiatives launched in recent years will gather much more. Nonetheless, much of the knowledge needed to build multi-level atlases and unifying models of the brain is still missing. Therefore, the goal of the HBP should be to generate and interpret strategically selected data that can act as a scaffold for future data generation, providing essential input for brain atlases and models. The HBP-PS has identified three main focuses for this kind of research.

Multi-level brain structure in mouse

Many of the basic principles governing the organisation of the mouse brain are common to all mammals. Furthermore, studies of mice can build on an enormous base of existing knowledge in genetics, molecular and cellular biology as well as on a large range of experimental techniques. However, discovering such principles requires systematic descriptions of the brains of genetically well-characterised animals at all possible levels of biological organisation. Such data sets do not currently exist. To meet this need, the HBP should systematically study the

relations between different levels of brain structure (gene sequences and chromatin structure, gene expression, protein expression, cells, synaptic connections, the neuro-vascular-glial system) in a cohort of mice expressing a variety of gene mutations. We recommend that the project test its methodologies in a detailed study of the mouse visual system, performed in collaboration with the Allen Institute, which has begun a large program in this area. Such a highly public effort would serve as a stimulus for on-going research by other groups working in the same field. Ultimately, we expect that most of the data needed to build unifying models of the brain would come from outside the HBP.

Multi-level structure of the human brain

Mouse data provides many insights into the human brain. Obviously, however, the human brain has unique features of its own. It is essential therefore to supplement the mouse data with parallel human data sets. Evident ethical considerations limit the methods we can use. Nonetheless, recent developments in non-invasive techniques provide new options for researchers. The project should use iPSC technology to generate data on gene and protein expression in different types of human brain cells, on the distribution of receptors on the cell surface, on cell morphologies, and on the numbers of neurons and glial cells in different brain regions. This data should be complemented with DTI data on the connectivity between different regions and with structural and functional MRI data on the shape and size of different regions at different ages. The data generated in this way would provide the essential scaffolding for the HBP atlas of the human brain (see below) and for model building. Additional data would come from predictive neuroinformatics. Easy access to this data through the HBP platforms would catalyse contributions from external groups, using them to gradually fill in the inevitable gaps in the initial data.

Brain function and cognitive architectures

To understand the “bridging laws” linking the physiological properties of neuronal circuitry to specific cognitive and behavioural competencies, we need to measure the dynamics of the brain as it performs well-characterised cognitive and behavioural tasks. These should cover a strategically selected range of human skills, from simple capabilities, also present in non-human species, to exclusively human skills such as language. To obtain this data, the HBP should combine human data from fMRI, DTI, EEG, MEG and other non-invasive techniques with behavioural data. The data from this work would help the project to develop high-level models of the cognitive architectures implicated in particular skills. Combined with performance data, such models would provide benchmarks for the validation of brain models and allow the construction of simplified models of the underlying neuronal circuitry, suitable for implementation in hardware. This would represent an essential first step towards systems with brain-like learning and information processing capabilities.



Figure 6: Precursor project: biologically detailed cortical microcircuits reconstructed and simulated by the Blue Brain Project, EPFL, Lausanne, Switzerland



Figure 7: Precursor project: ARM-based, many-core chips from the SpiN-Naker project, Manchester University, United Kingdom




Figure 8: Precursor project: wafer-scale integrated neuromorphic system from the BrainScaleS project, Heidelberg University, Germany



Figure 9: Precursor project: Europe's first Petascale computer. Forschungs Zentrum Jülich, Jülich, Germany



Figure 10: Precursor project: NeuroSpin - Brain imaging facility at CEA, Saclay, France

Theory

Today, theoretical neuroscience is highly fragmented and often has only a tenuous grounding in biological data. Yet without theory it will not be possible to effectively apply knowledge about the brain in medicine or computing. The HBP should thus include a cohesive programme of research addressing strategically selected themes essential to the goals of the project (see Figure 11): mathematical techniques to produce simplified models of complex brain structures and dynamics; rules linking learning and memory to synaptic plasticity; large-scale models creating a bridge between “high-level” behavioural and imaging data; and mathematical descriptions of neural computation at different levels of brain organisation.

We recommend that, as part of this programme, the project should establish a *European Institute for Theoretical Neuroscience*, encouraging participation by outside scientists and mathematicians, and acting as an incubator for approaches that challenge the conventional wisdom.

The mathematical principles underlying the relationships between different levels of brain organisation are essential to unite theories of how the brain represents information, learns and remembers

The integrative role of theory in the HBP

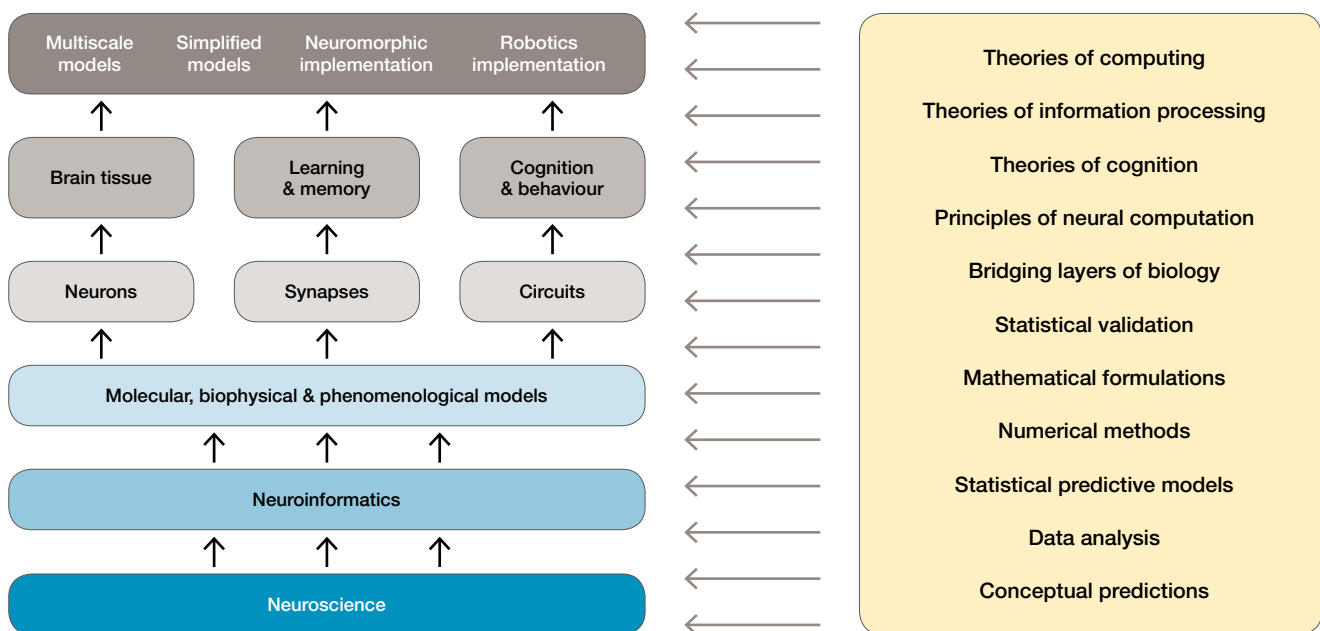


Figure 11: From biology to abstract mathematical representations

ICT platforms

The most important goal of the HBP should be to create a new technical foundation for brain research and its applications in medicine and computing. We therefore recommend that the HBP create an integrated system of six ICT platforms, dedicated to Neuroinformatics, Brain Simulation, Medical Informatics, High Performance Computing, Neuromorphic Computing, and Neurorobotics. If the European Commission accepts our recommendation, these platforms would become a valuable asset for European science and engineering, providing high quality, professionally managed services to these communities.

Neuroinformatics

The majority of current neuroscience is performed by relatively small research groups, each addressing well-focused research questions. Compared to other scientific communities (e.g. in physics or genomics) data sharing is rare and standards for measurement protocols, data formats and annotation are poorly developed. In these conditions, it is difficult to compare data and to integrate heterogeneous data sets in unifying models. In recent years, the INCF [3] and other international organisations (e.g. the Allen Institute [4]) have launched important initiatives to move beyond this situation. The HBP should collaborate with these organisations to develop new neuroinformatics tools for the analysis of brain structure and function and to encourage their widespread use. In particular the *Neuroinformatics Platform* should provide tools to manage, navigate and annotate spatially

*Predictive informatics
and brain simulation
make it possible to share,
predict, integrate
and unify data and knowledge
about the brain*

referenced brain atlases. The project should use these tools to create standardised multi-level atlases of the mouse and human brains, integrating data from the literature, other large-scale neuroscience initiatives, and smaller research groups, inside and outside the HBP. The publicly accessible atlases would provide new incentives for data sharing, and reduce unnecessary duplication of animal experiments. The HBP brain atlases would be the key source of data for the HBP's *Brain Simulation Platform* (see below).

A second key component of the platform should be tools for predictive neuroinformatics. Today, there are so many known gaps in the experimental data, and so many new gaps waiting to be discovered that it will take decades of experiments to build a complete picture of brain structure and function, even in a single species. To help overcome this problem, HBP tools for predictive neuroinformatics would detect statistical regularities in the relationships among different levels of biological organisation in different species, exploiting them to estimate parameters, which may not have to be measured experimentally. This technique has the potential to maximise the information extracted from experiments, reducing the number of duplicate experiments, and filling in gaps in our knowledge.

Brain Simulation

Understanding the brain requires a multi-level view of its structural and functional organisation. Today, however, neuroscience has no strategy for integrating the flood of data generated by experimental research. To overcome this fundamental obstacle, we recommend that the HBP should create a large-scale *Brain Simulation Platform*. The platform should make it possible to build and simulate unifying brain models that would integrate all the available data, generating “emergent” structures and behaviours that could not be predicted from smaller data sets. The platform would allow researchers to build models at different levels of description (abstract computational models, point neuron models, detailed cellular level models of neuronal circuitry, molecular level models of cells, synapses and small areas of the brain, multi-scale models that switch dynamically between different levels of description). This capability would allow experimentalists and theoreticians to build the models most appropriate to the questions they are seeking to answer. The platform would provide researchers with modelling tools, workflows and simulators allowing them to integrate very

Accelerated neuroscience

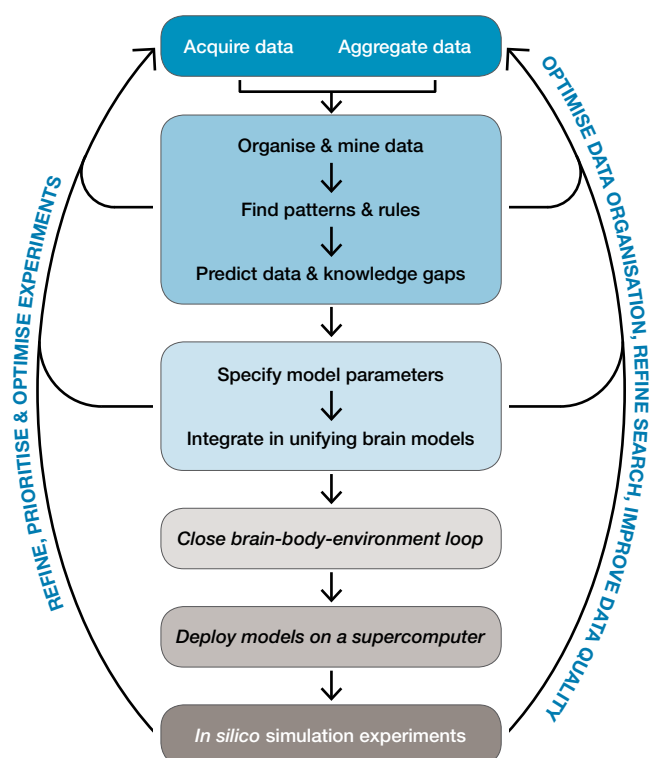


Figure 12: A positive feedback cycle for neuroscience: from data to simulation and back

large volumes of heterogeneous data in multi-scale models of the mouse and the human brains, and to simulate their dynamics. Researchers would use simulations as *virtual specimens*, in repeatable *in silico* experiments including systematic measurements and manipulations impossible in the lab (see Figure 12). Models developed in this way would play an essential role in HBP research in medicine (disease and drug simulation), neuromorphic computing (models of neural circuitry for incorporation in neuromorphic hardware) and neurorobotics (models to perform specific cognitive and behavioural tasks or to process specific classes of information).

High Performance Computing

Current supercomputing hardware lacks the power and software capabilities for multi-scale modelling of the human brain. And the cost and complexity of the technology prevent the majority of groups from using it in their research. Even more importantly, current paradigms provide no way of analysing and visualising the massive volumes of data that will be produced by exascale models and simulations¹. The HBP *High Performance Computing Platform* should offer a new solution, providing the project and the community with the supercomputing capabilities required

¹ In current paradigms, data is often stored externally for post-processing. With exabytes of data, this will become impossibly slow and expensive.

for multi-scale brain modelling, simulation and data analysis. These capabilities should be upgraded in stages, over the duration of the project, gradually moving towards the exascale. Simultaneously the project should develop completely new capabilities for interactive computing and visualisation, adapted to the needs of exascale simulations. These techniques would have immediate applications in the life sciences. In the longer term, they would allow the development of supercomputer-based services for industry, medicine and the consumer market.

Medical Informatics

An increasing body of literature demonstrates the feasibility of “mining” large volumes of clinical data for signatures of disease. However, nearly all current work treats diseases in isolation, focusing on comparisons between patients and controls. This approach makes it impossible to exploit the full range of variability present in large clinical data sets or to identify the unique biological characteristics of individual patients. The HBP should adopt a new approach, analysing data from very large numbers of patients and normal subjects, without pre-selecting patients with a particular diagnosis (see Figure 13). This strategy would make it possible to exploit resources rarely used for research, including massive hospital archives of imaging and other clinical data, and the proprietary databases of pharmaceutical companies.

Accelerated medicine

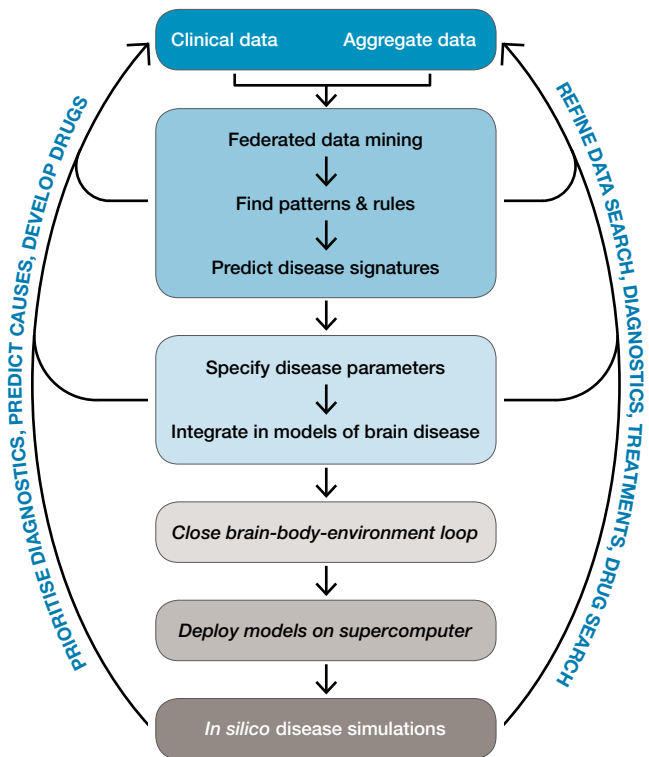


Figure 13: A positive feedback cycle for medicine: from clinical data to disease simulation and back

Identifying similarities and dissimilarities across brain diseases is a prerequisite for personalised medicine

Facilitated by the organisational structure of Europe’s largely socialised health services, the *Medical Informatics Platform* should federate this data, enabling researchers to query the data without moving it from the sites where it is stored and without compromising patient privacy. In this way, they would be able to study correlations and interactions between aetiology, phenomenology, nosology, diagnostic parameters, pathogenesis, treatment, and prognosis, identifying biological signatures associated with well-defined pathogenic processes. The identification of such signatures would enable them to develop new biologically grounded classifications of brain disease, leading to a new systematic understanding of its causes, and new diagnostic tools. These would not only facilitate early diagnosis – a precondition for effective treatment; they would also help to optimise the selection of patients for clinical trials, reducing costs and increasing success rates. Understanding the causes of disease would lead to new treatment strategies, new techniques for rational drug design, and new strategies for personalised medicine. The ultimate result would be improvements in treatment, better outcomes for patients, and reduced health care costs. The development of the necessary know-how would be led by European researchers. This knowledge base would help European pharmaceutical companies to lead the development of new therapies for brain disease – a segment of the drug market with enormous potential.

Neuromorphic Computing

According to the experts responsible for the International Technology Roadmap for Semiconductors [2], neuromorphic computing systems inspired by the architecture of the brain, are more likely than any other approach to overcome fundamental limits on the performance and capabilities of conventional computing technology. Developing such technologies requires advanced software (tools to simplify circuits for specific applications, design tools, tools for circuit testing etc.), supercomputing facilities (for circuit design and simulation), and hardware (neuromorphic devices and systems). Most importantly of all, it requires detailed knowledge of the neuronal circuitry and mechanisms responsible for specific cognitive functions. To meet this need, the HBP should create a *Neuromorphic Computing Platform* that provides researchers and application developers with capabilities to develop neuromorphic devices and systems with circuitry based on simplified versions of the neuronal circuitry of the brain, to test their capabilities, and to build prototype applications (e.g. neuromorphic systems for use in high performance computing, search and other advanced Internet services, industry, transport, and consumer electronics).

Once established as a working technology – the main goal of the platform – neuromorphic computing has the potential to become as important as the “conventional” computing technologies, which have fuelled economic growth for the last sixty years.

Neurorobotics

The brain evolved to control the body as it interacts with its environment, interactions that many researchers in robotics have attempted to replicate. In most cases, however, the models and robots they have used have been distant from biology. The HBP *Neurorobotics Platform* would allow them to interface a detailed brain model to a simulated body with an appropriate set of actuators and sensors, place the body in a simulated environment, train it to acquire a certain capability or set of capabilities, and compare its performance against results from human or animal experiments (see Figure 14). The platform would provide the tools and workflows necessary to perform this kind of experiment and to dissect the cognitive architectures involved, enabling radically new strategies for studying the multi-level mechanisms underlying behaviour. Additional tools for technology developers would make it possible to transfer brain models developed in simulation to physical robots, with low-power neuromorphic controllers – an essential step towards the development of robots – robotic vehicles etc. – for use in manufacturing, services and the home.

Accelerated future computing

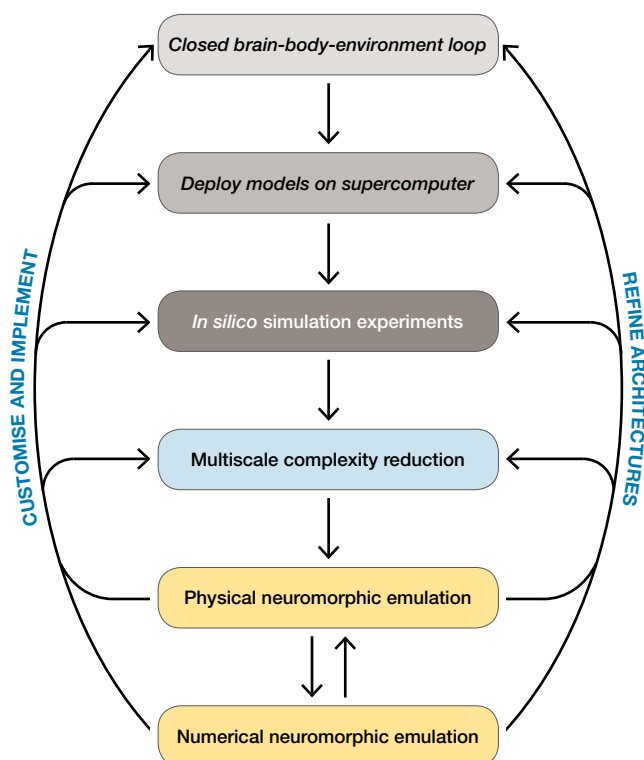


Figure 14: A positive feedback cycle for computing: from data to neuromorphic architectures and back

Applications

The value of the HBP depends on research and development enabled by the project’s ICT platforms. The project’s fourth goal should thus be to demonstrate how the platforms could be used to produce immediately valuable outputs for neuroscience, medicine and computing. Initial pilot projects should be managed by HBP partners with the necessary background and know-how. Pre-competitive research and development should be organised as collaborations with industry. However, the majority of this work should take the form of projects proposed by groups and researchers from outside the initial HBP Consortium, chosen in response to competitive calls for proposals (see page 70).

We recommend that the HBP work programme should focus on three priorities: integrative principles of cognition; understanding, diagnosing and treating brain disease; and future computing technology.

Integrative principles of cognition

Neuroscience and medicine both require an integrated multi-level understanding of brain function in the context of cognition and behaviour. The HBP ICT platforms would provide a new foundation for this kind of research. Once brain models have been integrated with a simulated body acting in a simulated environment and trained to display a particular competency, neuroscientists would be able to systematically dissect the neuronal mechanisms responsible, making systematic manipulations and measurements that would be impossible in the lab. Similarly, medical researchers would be able to simulate the effects of alterations at different levels of brain organisation (gene knockouts, knockouts of specific types of neurons, changes in topology) and model the effects of drugs and other treatments. Pilot projects would use the capabilities of the HBP to investigate cognitive capabilities investigated in HBP work on brain function and neuronal architectures (perception and action, decision-making, goal-oriented behaviour, navigation, multisensory perception, object recognition, body perception). HBP calls for proposals from outside groups and researchers should support new work on these problems and in other areas of cognitive science, proposed by researchers themselves.

Understanding, diagnosing and treating brain disease

Today, medical researchers lack the data and the tools to understand the causes of brain disease and to develop new treatments. The *Medical Informatics Platform* would allow them to study disease mechanisms systematically, analysing every possible category of clinical data (genes, gene expression and other “-omics” data, imaging features etc.) from healthy subjects and from patients with widely varying conditions. The ability to analyse very large, clinically diverse populations would enable them to identify biological signatures (groups of features in the clinical data) associated with specific kinds

Demonstrate the value of the HBP platforms, making them a vital asset for the scientific community

of lesion and specific defects of cognition and behaviour. The discovery of such signatures would lead to hypotheses about their underlying causes that the *Brain Simulation Platform* could model and test. A better understanding of the biological causes of brain disease would lead to new biologically grounded systems of classification, and better diagnosis. It would also lead to new strategies for treatment and new drug targets. The *Brain Simulation Platform* would make it possible to test these ideas before embarking on animal experiments or clinical trials – avoiding the cost of testing drugs with poor prospects of success – and speeding up the drug discovery process.

Initial research should focus on high impact diseases such as autism, depression and Alzheimer’s disease. From the start, however, the project should adopt a systematic approach, addressing the multi-level biological mechanisms responsible for the symptoms, prognosis and response to treatment of individual patients. As the project progresses, an increasing proportion of research should be performed by researchers from outside the initial Consortium, selected via open calls for proposals. Researchers would be free to propose their own topics for research. Other work should be developed in collaboration with the pharmaceutical industry. The end result should be new diagnostic tools allowing early treatment of diseases before they cause irreparable damage, new techniques of personalised medicine, new tools for drug discovery and screening and ultimately new treatments. The first results are likely to come in the HBP’s first few years. Given the huge impact of brain disease on European health services and European citizens even small improvements would bring large benefits.

Future computing technologies

Conventional computing is rapidly approaching fundamental limits on processing speed, power consumption, reliability and programmability. This has led to proposals to complement current ICT with computing systems, architectures and software inspired by the architecture of the brain. To date, however, relatively few research groups have had access to the necessary technologies and knowledge, and none have had access to biologically accurate models of neuronal circuitry. The *High Performance Computing*, *Neuromorphic Computing* and *Neurorobotics* platforms should meet this need.

The *High Performance Computing Platform* should allow researchers to explore hybrid technology integrating neuromorphic devices with “conventional” supercomputing for specific classes of computation. Other interesting themes include brain-inspired communications protocols, and brain-inspired storage and retrieval.

The *Neuromorphic Computing Platform* should enable the development of prototype systems for prediction making, data mining, spatial and temporal pattern detection. Some of this work would involve the development of special-purpose neuromorphic chips. Partnerships with industry should explore the development of neuromorphic controllers for vehicles, household appliances, manufacturing, image and video processing, and telecommunications.

In neurorobotics, the HBP ICT platforms should provide researchers with standardised workflows for the development of new applications, including potentially valuable neuromorphic and robotic systems, and custom robots for specific applications (e.g. industrial and service robots,

Energy scales



Figure 15: From biological brains to supercomputer simulations. Energy per synaptic transmission spans fourteen orders of magnitude. Neuromorphic computing occupies an intermediate position between biological brains and supercomputer simulations

automatic vehicles, medical robots, robots for ambient assisted living etc.).

As in the case of medicine, the potential impact is very large. Neuromorphic computing will not replace the computing technologies we know today. Rather it could take on a complementary role, offering services requiring forms of brain-like intelligence and flexibility that are difficult or impossible to implement on current systems. The market for services using the new technologies could be as large as the market for the current generation of ICT.

Why a FET Flagship?

The Human Brain Project is an ideal fit for the goals of the FET Flagship Programme, for eight reasons.

1. The HBP would address scientific challenges with a very large potential impact. Understanding the chains of causation leading from genes, cells and networks to cognition and behaviour would be a major achievement for neuroscience; unravelling the biological mechanisms of brain disease would bring huge benefits for European health budgets and for European citizens; brain-inspired computing technologies have the potential to transform European industry and society.
2. The project has a well-defined goal and well-defined plans for data generation, platform building and applications development. The project's work plan provides a clear focus for research and clear success criteria.
3. The HBP cannot realise its vision without a long-term programme of research involving a large number of partners and significant investment of financial and technical resources. Such a programme would not be feasible without a project on the scale of a Flagship.
4. The project offers an unprecedented opportunity for interdisciplinary collaboration across a broad spectrum of fields, from neuroscience to high performance computing. No single country has more than a small part of the necessary know-how and resources. European scale collaboration is essential for the project's success.
5. The project, while very challenging, is technically feasible. The techniques it would deploy and the capabilities it would build are firmly rooted in existing science and technology and in technologies (such as exascale computing) that already have a well established development roadmap. In simulation and modelling, the project would build on EPFL's Blue Brain Project, which has already prototyped many of the necessary modelling tools and workflows; the project's effort in high performance computing would be led by FZJ, which also leads PRACE, Europe's largest supercomputing initiative. In medicine, HBP activities would be led by CHUV, which would coordinate a consortium of major European hospitals. HBP work in neuromorphic computing and neurorobotics would build on work in other projects in the FET programme (FACETS [5] BrainScales [6], ECHORD [7] and in nationally funded projects (SpiNNaker [8]). The leaders of these projects have agreed to participate in the HBP should it be approved.
6. The project has the potential to produce economically and socially valuable results long before it achieves its ultimate goals. All the planned technology platforms would be operational within the first two and a half years. The first informatics-based techniques for the diagnosis of Alzheimer's and other neurological and psychiatric disorders should be available within the first 3-5 years. The HBP would also play a critical role in the development of high performance computing, both for general applications and for the life sciences. One key development would be the development of novel techniques of interactive supercomputing – allowing the construction of “virtual instruments” (e.g. virtual MRI machines) equivalent to the physical instruments used in biological and clinical research. On the same timescale, the project would develop the first neuromorphic computing systems based on simplified, yet biologically realistic versions of brain circuitry, and the first neuro-robotic systems for use in closed-loop experiments and applications development.
7. The HBP's goals closely match European and national research priorities. All European countries emphasise human health as a key priority for research, often placing a special emphasis on aging. The HBP's emphasis on the causes, diagnosis and treatment of brain disease matches these goals. The project's work on neuromorphic computing can make an important contribution to a second key priority: energy and the environment. Neuromorphic computing technology offers the promise of a new category of computing systems, with a range of possible applications as large as current systems but with far lower energy consumption. The third key priority for all European governments is European industry's ability to compete on world markets. Particularly critical is the situation in the pharmaceutical industry, which is cutting back its Eur 4 billion annual research budget for brain disease. HBP research on the causes and classification of brain disease, together with simulation-based pharmacology can help to make this research profitable again, allowing European industry to take a leading role. HBP-led innovation in high performance computing, neuromorphic computing and neurorobotics can establish European leadership in future computing technologies of potentially vital importance for the world economy of the 2020s and 2030s.
8. The HBP perfectly matches an emerging priority, not yet fully visible in national and European research programmes: namely the urgent need for massive investment in brain research. Public awareness about the aging of the population and the burden of neurodegenerative disease is already high. The recent EBC reports on the costs of brain disease [1] has rung new alarm bells. There is a growing consensus that Europe needs a large-scale research effort to understand the brain and its diseases. The HBP would meet this need.

3 Science and technology plan

Data

Multi-level structure of the mouse brain

Objectives

To understand the role of different levels of biological organisation in human brain function, it is essential to understand how they work together in other mammals (see Figure 16). For this we need systematic data describing different levels of brain organisation in a single species and analysing how variation in structure relates to genetic variation. Much of our current knowledge of the genetics, molecular and cellular biology and behaviour of the brain and many of our experimental techniques come from studies in mice. We therefore propose that the Human Brain Project should generate systematic mouse data sets for genomes and molecules, cells and circuits, using the results to fill in gaps in our current knowledge and to discover general principles we can apply in multi-level models of the human brain.

Generate a skeleton of strategically selected data as a catalyst for community research

State of the art

Current neuroscience comprises a vast range of different disciplines and research communities, each focusing on a specific level of biological organisation, and on the brain regions, species and methods best adapted to its specific goals. Progress is rapid at all levels. However, data and knowledge are badly fragmented. At the cellular anatomy and connectivity levels, we still lack complete data for any species. Even in *C. elegans* – the only species whose neuronal circuitry has been completely deciphered – we are still missing essential information, such as data on neural morphologies. At the physiological level, we do not have a clear, quantitatively accurate picture of physiological response in different types of synapse, cell and circuit. Data on long-range connections between different brain regions is similarly sparse. In the absence of data, we cannot understand the relationships between different levels of brain organisation – for instance, the way in which a variant in a specific gene can affect the architecture of an animal's neural circuitry and its subsequent behaviour.

Today, however, new technologies are rapidly making it easier to generate the data we need, and to relate data in mouse to data in humans. At the molecular level, we already have a large volume of quantitative data, including data on DNA sequence and modifications [9], RNA [10] and proteins [11, 12]. Work is in progress to build molecular-level atlases of the human and mouse brains. Combined with RNA and protein profiles for different cell and synapse types, these atlases would soon make it possible to estimate the numbers of cells of different types in different brain regions and to relate the data for the two species.

At higher levels of organisation, breakthroughs in scalable methods, and particularly in optogenetics [13] are paving the way for comprehensive studies comparable to the work currently in progress in molecular biology and proteomics. There has also been considerable progress in connectomics. Molecular tracer methods make it possible to trace connections between different types of cells and their synapses (using molecular and structural approaches). Data from these studies can be correlated with results from behavioural studies, traditionally performed in low-throughput settings, but now complemented by high-throughput methods using touchscreen based perception and learning tasks [14]. These methods mean it is now possible to measure and compare data from thousands of animal and human subjects.

Methodology

We propose to base the HBP's work in this area on a cohort of 200 mouse strains expressing a range of sequence variants (mutations and normal gene variants). This method would allow the project to systematically generate strategically valuable data (DNA sequences, chromatin, mRNA and protein expression, synaptic connections and cell structure, physiology) and to compare the results against human data sets (see Figure 17). The results would allow the project to elucidate the causal relationships between different levels of brain structure. Organisational principles derived from this work would help the HBP to estimate parameter values for features of the human brain that cannot be measured experimentally.

The study should seek to answer key questions concerning the relationships between different levels of brain organisation.

1. *The genome.* What is the relationship between gene sequences and chromatin modification and higher levels of brain organisation (gene and protein expression,



Multi-level data required for multiscale brain models

Cognition. Structured cognitive tests, combined with fMRI and MEG, constrain brain circuit models for visual perception-action, decision and reward, memory encoding and retrieval, space, time, number and language. Changes in circuits during development and aging and differences between adults constrain network models of cognition at different ages and provide benchmark data to validate detailed cellular brain models.

Whole brain. Multi-modal sensory physiology and anatomy map how different brain regions come together to shape perception-action and our sense of body ownership, awareness and consciousness, and can help validate human brain models. Whole brain synchrotron scans reveal the vasculature supporting cognition, and provide constraints for models of blood flow and brain metabolism.

Connectivity. Whole brain fibre tract tracing and DTI yield paths and fibre densities for connectivity within and between brain regions and provide global fibre projection parameters constraining connectivity in brain models. Whole brain ultramicroscopy yields constraints for single neuron projections and provides high-resolution validation of DTI tracts. EM provides high-resolution images of the synapses formed on individual neurons, principles of fibre selection, and structural features of synapses.

Brain regions. Structural and functional MRI yield dimensions of brain regions which can be used to build models. Region-specific cellular architecture and densities constrain the cellular composition of model regions. Receptor, ion channel, signalling and other protein distributions further constrain neurochemical organisation within and across brain regions. Correlations between protein distributions, cognitive circuits and genomic variability point to neural mechanisms of cognition, provide global constraints for detailed brain models and generate data for model validation.

Microcircuits. The cellular and molecular composition of microcircuits supports their role in cognition. Single-cell gene expression yields sets of genes that form different types of neurons and glia and determine their morphological and electrical properties. Global brain maps of gene and protein distributions constrain the cellular composition of microcircuit models. Cell geometry and synaptic selection rules constrain local synaptic connectivity in microcircuit models. Electrophysiology, multi-electrode recordings, voltage sensitive dye mapping and optogenetic studies provide data to validate microcircuit models.

Cells. 3D reconstruction of the anatomy of single cells yields the structural geometry needed to establish the morphological properties of different cell types. Correlations between gene expression and the geometric properties of cells constrain the artificial synthesis of cellular morphologies from gene expression patterns, as well as models of morphological plasticity. Single-cell gene expression, combined with general rules for the production and distribution of proteins and for protein interactions, constrain molecularly detailed models of neurons and glia.

Synapses. Physiological, biophysical and imaging studies of synaptic transmission, plasticity and neuromodulation constrain synaptic models. Pair-wise single cell gene expression constrains the repertoire of synaptic proteins at the synapses between pairs of neurons of known type, making it possible to model synapse diversity. The dynamics of single cell gene expression constrain long-term molecular changes in synapses in response to environmental stimuli. Comparing synaptic proteins across species constrains species-specific synaptic models.

Metabolome. The intricate biochemical network linking neurons, synapses and glia constrains molecular brain models, in which activity, plasticity, neuromodulation, homeostasis and nutrition depend on metabolic processes. Coupling neurotransmitter receptors and their signalling pathways to the biochemical pathways that supply energy to cells and synapses, constrains activity-driven changes in blood flow, and the resulting fMRI signals.

Proteome. The number and different types of proteins cells produce, the different parts of the cell where they are located, and their respective life cycles all constrain how many and which proteins can come together in a single cell. The set of proteins, other biochemicals, and ions that each protein binds to and reacts with forms the protein's interactome. Results on protein-protein interactions from biochemical studies, molecular dynamic simulations, and predictive informatics constrain reaction-diffusion models of cells.

Transcriptome. Combined with data on the genes expressed in single cells, 3D maps of gene expression constrain the total number of neurons in the brain and the types of genetically identifiable cells the brain can produce. Single cell gene expression changes in response to stimulation, determining how cells can change with experience. Single cell gene expression, combined with predictions of which proteins the cell can produce and basic principles of proteomics, constrains detailed molecular-level cell models.

Genome. The state of the chromosome reflects when genes are active or inactive and constrains gene network models. It is likely that the genome constrains the number of genetic cell types in the brain, the size of brain regions, connectivity between brain regions, and total brain size. It may also predict cognitive functions, behavioural traits, epigenetic vulnerability, and brain disorders.

Figure 16: Generating a skeleton of multi-level data as a catalyst for future community research

Multi-level structure of the mouse brain




Figure 17: Roadmap for the generation of mouse structural data (MSD). Generation of physiological data will be entrusted to research groups from outside the initial HBP Consortium, selected via competitive calls. The data will provide vital input for unifying models of the mouse brain

- distributions and densities of different cell types, connectivity, the size of different brain regions, large-scale structure of the brain)? How can we characterise the cascade of multi-level effects leading from genes to behaviour?
- Gene expression.** What combinations of genes are expressed in different types of cells at different ages? What are the dynamics of gene expression in single cells? What can we predict about the cell, by reading mRNA profiles? What are the mechanisms underlying spontaneous, stimulus and environmentally driven changes in gene expression?
- Protein expression.** What proteins are expressed in different types of neuron, glia and synapse? How does protein expression affect the electrical and pharmacological behaviour of cells? What are the molecular and cell biological principles governing the distribution of proteins within the cell? What can we learn from these distributions?
- Cells.** How many and what types of cells are present in different regions of the brain? What are their morphologies? What are the relationships between genetic mutations, gene expression and morphology?
- Connectivity.** How many different types of synapse are there? How do neurons select their targets and choose their synaptic locations? How do brain regions map onto each other? How many brain regions can a neuron project to?
- Electrophysiology.** What are the different profiles of excitability and firing patterns of different neuronal types? What are the different types of synaptic plasticity? What are the mechanisms underlying diversity in synaptic transmission? How do neurons process synaptic input? What are the characteristic emergent behaviours of neuronal microcircuits? How do microcircuits of neurons work together to shape the dynamics of a brain region?
- The neuro-vascular-glial system.** How do neurons, glia and blood vessels interact? What is the detailed architecture of the vasculature that directs blood within the brain? What is the structural relationship between neurons, glia and vessels? How do changes in neurons alter the properties of vessels and *vice versa*?
Combined with behavioural data from elsewhere in the project, this data would provide fundamental insights into the way brain regions interact to shape perceptions and support cognition. These insights can help the project to answer new questions. Which combination and sequence of activation of different brain regions support different forms of behaviour? How do genes and gene expression correlate with cognition and behaviour? How are the building blocks of behaviour related to one another and what is their mechanistic underpinning at the molecular, cellular and circuit levels? What is the smallest network of neurons that can perform an isolated task? How does the composition of a neural microcircuit affect the computational operations it performs?

What is the role of single cell types in the processing of sensory and motor information? How important is multisensory information processing for the individual senses?

HBP data generation would provide only a very small fraction of what is needed to build detailed models of the brain. Nonetheless it would create a basic skeleton that could then be fleshed out with data from other groups and from predictive neuroinformatics (see below). To reinforce this approach, we propose that the HBP collaborate with the Allen Institute in their on-going multi-level case study of the mouse visual system. At the structural level, the study would map the volumes of brain areas implicated in the visual system, obtaining cell numbers and distributions, as well as data on genetically characterised cell types and on neuron morphologies. Functionally, it would identify the role of single neurons and cell types in visual information processing and visual perception and learning paradigms. The results generated by this work would be contributed to the INCF (www.incf.org), the GeneNetwork system (www.genenetwork.org), the Genes to Cognition programme (www.genes2cognition.org) and other internationally established databases, which would provide standardised data resources for theoretical studies and for modelling. The Allen Institute has confirmed that it would be willing to participate in such a study.

Multi-level structure of the human brain

Objective

Mouse data provides many insights into the human brain. Obviously, however, the human brain is different. It is essential therefore to supplement mouse data with direct human measurements. Although ethical considerations limit the methods we can use, recent developments in non-invasive techniques provide new options. We propose that the HBP use these techniques to generate a scaffold of strategically selected data on the structure and functional organisation of the human brain at different ages and at different levels of biological organisation (see Figure 18). It would then use this scaffold to catalyse and organise contributions from outside the project, filling in the gaps with data from predictive neuroinformatics (see p. 33). The results would provide essential input for multi-level models of the human brain and for the understanding of human brain disease.

State of the art

Genetics and gene sequencing. Genetics is the method of choice for understanding genome-to-phenome linkage at the molecular, cellular and behavioural levels. Two genetic strategies have proven particularly valuable. The first compares the phenotypes produced by point mutations against controls; the second examines small populations of individuals and assesses the role of endogenous genetic variation (natural polymorphisms). Combined with massive “-omic” data sets, these approaches make it possible to build and test complex systems models where every trait, at every level and scale, can be linked back to a set of gene loci [15]. The recent introduction of computerised touchscreen approaches

has made it possible to compare a subset of human cognitive functions with equivalent functions in mouse [16]. Despite the limitations of mouse models for predicting complex behaviour and cognition in humans, comparative studies of mice and humans can provide valuable information about putative mechanisms. Functions amenable to this approach include attentional processing, forms of associative learning such as visual and auditory memory, as well as cognitive flexibility and response inhibition. These methods provide a valuable tool for studies of normal human genetic variation.

Human mutations are a major cause of brain disease. Studies have identified over 200 single gene mutations affecting human postsynaptic proteins and over 130 brain diseases in which these mutations are believed to play a role. Studies of individuals with these mutations can provide useful insights into the way variation in specific proteins contributes to differences in cognitive, behavioural and emotional phenotypes while simultaneously providing valuable information on mechanisms of disease causation. Particularly interesting in this respect are studies on gene-carriers, with no overt signs of disease.

Molecular systems biology. Molecular systems biology uses mathematical and computational methods to understand the molecular basis of information processing in the brain. For example, multi-scalar analysis of genomic variation data and quantitative phenotype data make it possible to map patterns of gene and protein expression to specific neuronal and synapse types. With massive, well-structured molecular data for key brain cell and synapse types, it becomes possible to build rich quantitative models of higher order components – synapses, cells, neuronal ensembles and brain areas and to link these models to precisely matched anatomical, functional, and behavioural data sets, a precondition for predictive modelling.

Cataloguing cell types using transcriptomic data. Large-scale mapping of gene expression patterns in the mouse brain [17, 18] has confirmed that morphologically distinct cells express different combinations of the same genes. The Allen Institute is now conducting similar studies on human brain material from biopsies and *post mortem* examinations [19]. Combined with data from single cell transcriptomics – not yet available but on the horizon – this data would make it possible to predict the cell types present in different regions of the brain. In principle, the data could also enable prediction of the proteins present in different types of cells.

Cataloguing synapse types using proteomic data. Proteomics studies of human synapses have demonstrated that human synapses contain over a thousand different proteins [3]. The results indicate that the protein composition of synapses differs between different brain regions, different neuronal types and even along the same dendrite, and that certain patterns of synaptic protein are typical of specific cell types and brain regions [20]. Array Tomography, a new technique, makes it possible to analyse between ten and

Multi-level structure of the human brain

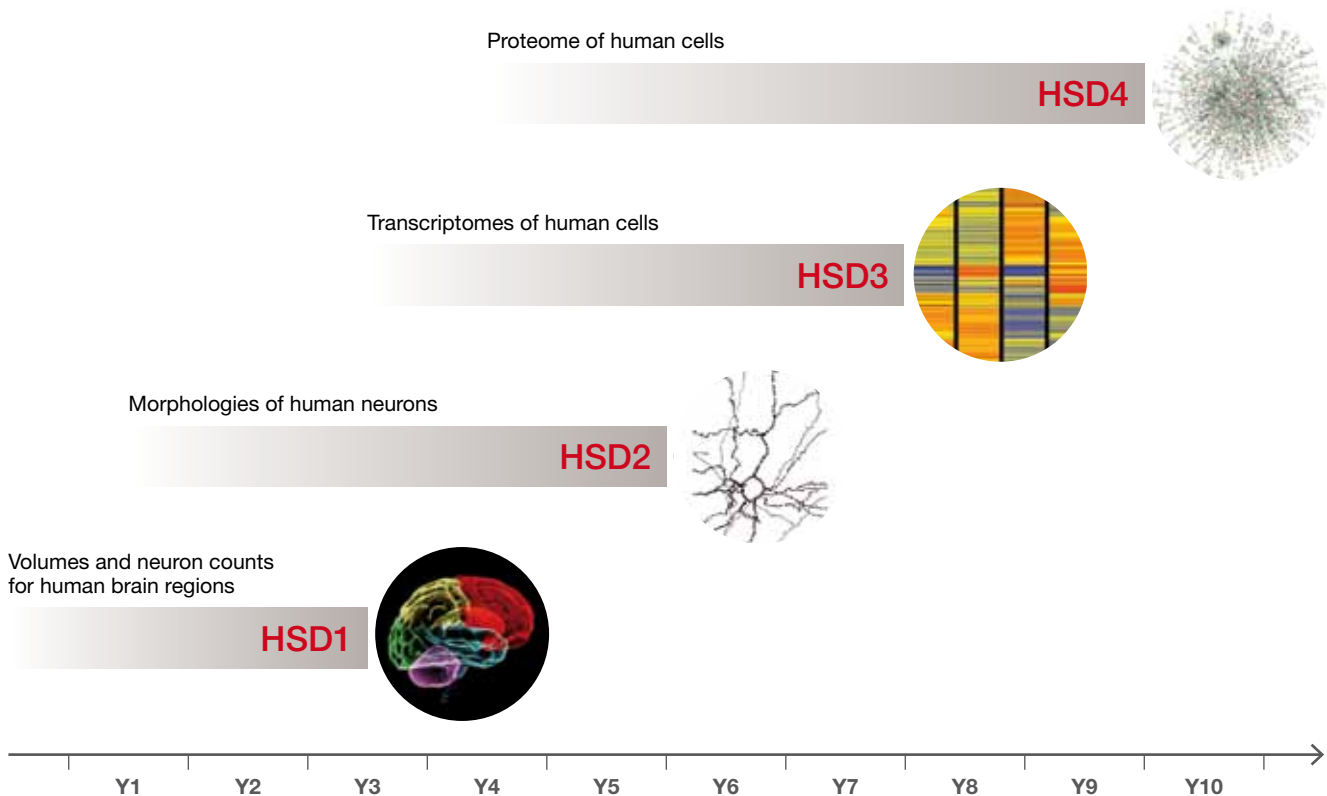


Figure 18: Roadmap for the generation of human structural data (HSD). As far as possible, the data collected should match the mouse data. It will be used to build and validate unifying models of the human brain

twenty synaptic proteins, mapping synapse diversity at the single synapse level [21]. Recently developed optogenetic methods for labelling synaptic proteins allow rapid, highly efficient mapping of individual synapse types, characterisation of the synapses present in different regions of the brain and identification of their role in neuronal information processing.

Living human neurons from stem cells. It is now possible to study living human neurons derived from human induced Pluripotent Stem Cells (iPSCs) [22]. The combination of iPSCs with developmental neurobiology has made it possible to generate a sufficient diversity of neurons and glia required to model human cortical function in a dish [23]. In particular, the zinc finger nuclease technique makes it possible to generate human neurons carrying disease mutations. These can be used to model neurodegenerative and psychiatric disease [24].

Imaging. Structural and functional imaging of the living human brain provide a valuable supplement to high-resolution data from anatomical *post mortem* studies [3]. Maps of the density of the main types of neurons obtained in *post mortem* brains provide a link between functional imaging data and underlying brain anatomy [20]. More recently, *in vivo* imaging techniques, particularly diffusion imaging and resting state imaging, have made it possible to begin mapping

the human connectome [21-23]. Polarised Light Imaging (PLI), detecting the myelin surrounding axons, makes it possible to link this data to the microscopic level and to verify *in vivo* data [24]. Intra- and subcortical connection profiles for individual areas, obtained in this way, are likely to provide new insights into the structure and function of the brain. For the human brain, PLI is also one of the few methods that can bridge the gap between the macroscopic organisation of the brain and knowledge about long and short fibre tracts, including those within the cerebral cortex. Given that most current information on human brain connectivity is extrapolated from animal and developmental studies, this is a crucial step.

Post mortem studies provide useful information about the distribution of different types of transmitter receptor in different regions of the brain [25]. Receptors play a key role in neurotransmission and are highly relevant for understanding neurological and psychiatric diseases and the effect of drugs. Diverse sets of molecular and phenotype data including disease data, brain physiology and behavioural data have been used in simulations of synapse function [26]. So far, however, most of this work has been based on static interaction representations that do not capture the full molecular dynamics of the nervous system. Molecular dynamics models would require HBP high performance computing capabilities.

Cellular brain models should be constrained by precise data on the cellular organisation of different brain areas and their connectivity

Recent evidence suggests that many neurological and psychiatric diseases (e.g., epilepsy, schizophrenia, major depression) depend not so much on single receptors as on the equilibrium among multiple receptors. Modelling and simulation provide an essential tool for understanding these complex mechanisms.

Brain models require precise data on the cellular organisation of different brain areas (e.g. cortical layers and columns) and their connectivity. Recent studies have combined *post mortem* studies of laminar cell distributions with *in vivo* diffusion techniques to measure the distribution of cell and fibre diameters, opening the road to *in vivo* studies of human cytoarchitecture and connectivity.

Methodology

The single cell transcriptome. The HBP should measure the single cell transcriptome of specific types of clinically derived brain cells and when this becomes possible, from human iPSCs. It should then compare the data with data from mouse studies. Combined with gene expression maps and modelling, this data would make it possible to predict many aspects of brain structure that cannot be measured experimentally.

The proteome. The HBP should measure the proteins expressed in human neurons, glial cells and synapses, and compare the results against data from mice.

Distribution of receptors. This work should map the distribution of receptor types of different neurotransmitters in different brain regions. The results would provide a solid basis for modelling neurotransmission, neuromodulation and the effects of psychoactive drugs and toxins.

Neuron morphologies. This study should characterise the morphologies of different types of neuron present in different regions of the human brain. Combined with modelling, the results would enable the project to predict a large proportion of the short-range connectivity between neurons, without measuring the connectivity experimentally.

Neuronal architecture. Neuronal architecture differs between brain regions with respect to the density, size, and laminar distribution of cells, and the presence of cell clusters. Significant differences have been observed in primary vs. secondary, visual vs. auditory, sensory vs. motor, and phylogenetically old vs. younger areas. This work would map the architectures of different layers and areas of the brain, providing constraints for simulation and modelling by introducing area-specific information on the level of large cognitive systems and behaviour.

Human brain connectomics. The HBP should use Diffusion Tensor Imaging (DTI) and Polarised Light Imaging (PLI) to derive patterns of connectivity between brain regions and to identify fibre tracts connecting layers and cells within brain regions. This data is essential for modelling the large-scale structural architecture of the brain.

Mapping of the developing, adult and aging brain. Structural and functional MRI would make it possible to map inter-individual differences in the adult human brain, and to identify structural changes characteristic of different stages of development and aging. Such information is necessary, among other reasons, to understand and model the formation of fibre tracts, the development of human cognition and the transition to disease.

Brain function and cognitive architectures

Objective

The goal of this work should be to take a set of well-defined cognitive tasks, already partially studied by cognitive neuroscience, and to dissect the organisation of brain activation and response dynamics as the brain performs the tasks. These studies should span scales ranging from global networks to local cortical maps and, where possible, sets of individual neurons (see Figure 19). The resulting data would allow the project to develop high-level models of the cognitive architectures implicated in particular competencies. Combined with behavioural performance data, they would provide benchmarks for the validation of the detailed brain models produced by the *Brain Simulation Platform* and guide the development of simplified models for use in neuromorphic devices and neurobotic systems.

State of the art

Functional specialisation of the human brain. While the first demonstration of a link between brain structure and cognitive function came from *post mortem* neuro-anatomy, the recent neuro-imaging revolution has greatly refined our understanding of cortical and subcortical functional specialisation [27]. Thanks to these techniques, we now have relatively precise information about the areas of the human brain responsible for processing particular categories of visual information (e.g. information on faces, body parts, words), for so-called core knowledge systems (systems handling information about space, time or number), for language processing, and for representing other people's minds (theory of mind).

Neural codes. The localisation of the areas responsible for specific functions is a means, not an end. Recent studies have attempted to characterise areas and regions in functional terms, i.e. to study how activation varies with stimuli and tasks, and to understand internal coding principles. High-resolution fMRI, repetition suppression and multivariate analyses of activation patterns form an essential toolkit that, in the best cases, allows precise inferences about the underlying neuronal codes [28, 29].

Function of the human brain

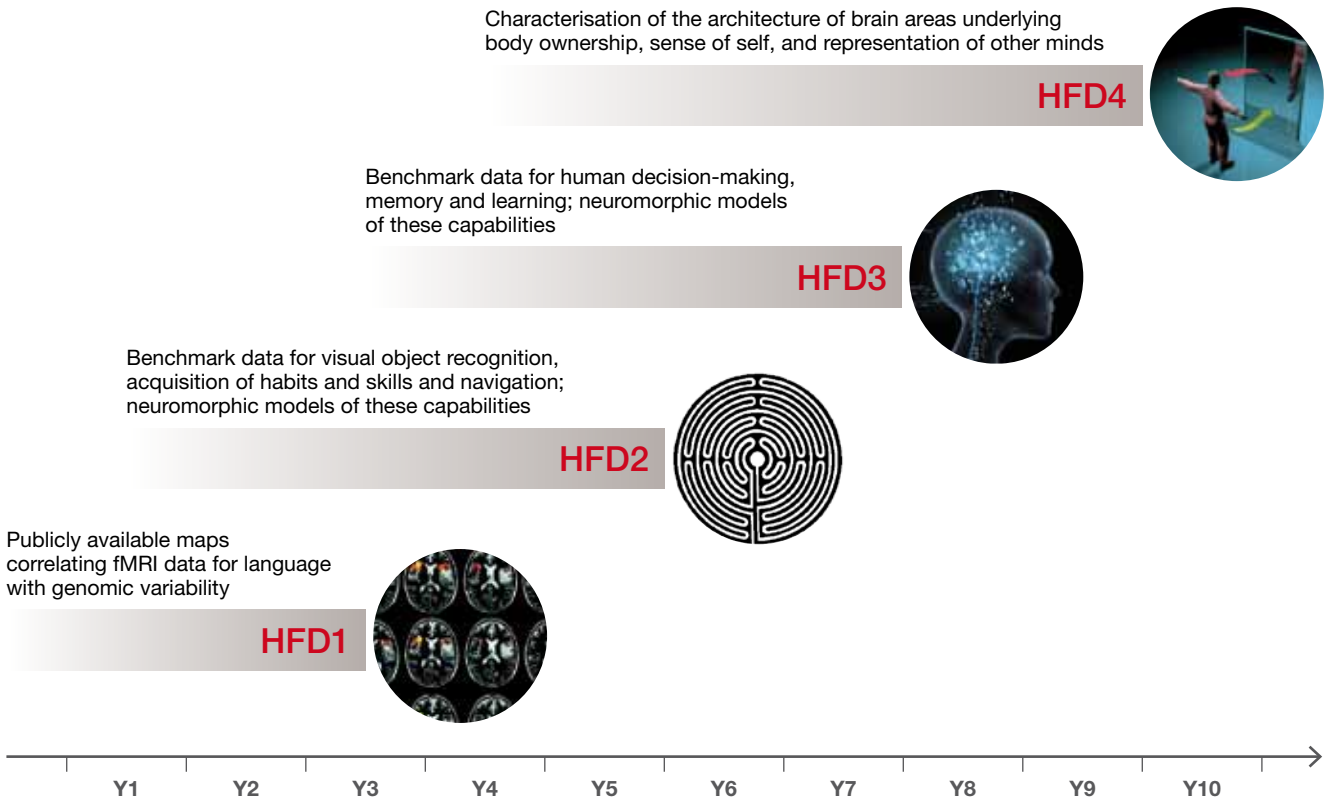


Figure 19: Roadmap for the generation of human functional data (HFD). The data will be used to develop theories of brain function, to build conceptual brain models and as benchmark data for the validation of human brain simulations

Spontaneous activity. Further insights come from studies of the way functional activity changes over time, including “resting state” studies, in which brain activity fluctuates “spontaneously” [30]. While some scientists see these fluctuations as nothing more than a consequence of neural noise in a non-random structural network, others interpret them as a dynamic internal model of the environment [31]. What is certain is that continuous spontaneous activity is a key characteristic of the brain that distinguishes it from engineered information processing systems. Understanding resting states and their dynamics could provide a strategy for systematically parsing functional brain areas and circuits in the living human brain.

Neurophysiological dynamics in human and non-human animals. Timing information from fMRI has made it possible to parse the dynamics of language and executive networks at ~200 millisecond resolution [32, 33]. A greater level of spatio-temporal detail on local codes and their dynamics is provided by electrophysiological recordings, using non-invasive MEG and EEG in humans, intracranial grids and single-electrode recordings in epilepsy patients, and grids and multi-electrodes in non-human primates. Neural codes have been identified for high-level vision and decision-making, including human cells responsive to objects, faces, places, and people. Another example is the study of attention. A prominent proposal suggests that attentional filtering

is implemented by selective synchronisation among neurons representing behaviourally relevant information [26]. This proposal can be tested by multiple neuronal recordings [34].

High-level cognitive functions. The recent literature includes descriptions of networks for language comprehension, reading, and mathematics, and the way they develop from infancy to adulthood. Other studies have focused on the way humans and other primates form strategies, detect errors and switch between tasks. This work has shown how networks crossing the prefrontal and parietal regions implement a “central executive”, a “global neuronal workspace” and a “multiple-demand” system [35].

Capabilities unique to the human brain. A major question currently being addressed by comparative studies concerns which cognitive abilities, if any, are unique to humans [36, 37]. These studies show that, at the sensory-motor level, humans and other primates are highly similar in many respects [38, 39]. Humans are distinguished by their recursive combinatorial ability, the capacity to bind words or other mental objects into hierarchical nested or constituent structures, as seen in the formation of linguistic sentences, music or mathematics. Recent studies have identified neuronal networks associated with these capabilities (see for example [40, 41]). Monkeys can perform elementary arithmetic operations similarly to humans, and even acquire symbols for digits [42]

but only humans seem able to “chain” several operations into nested algorithms [37]. Finally, the human brain may have a unique ability to represent an individual’s own mind (second order or “meta” cognition) and the thoughts of others (“theory of mind”). fMRI studies have identified a reproducible social brain network, active during theory of mind, but also during self-oriented reflections and the resting state. Interestingly, this network is modified in autism [43, 44].

Methodology

The HBP should combine human data from fMRI, DTI, EEG and other non-invasive techniques, identifying and characterising the neuronal circuits implicated in specific well-characterised cognitive tasks. The project should focus on the following functions.

- *Perception-action*: invariant visual recognition; mapping of perceptions to actions; multisensory perception of the body and the sense of self.
- *Motivation, decision and reward*: decision-making; estimating confidence in decision and error correction; motivation, emotions and reward; goal-oriented behaviour.
- *Learning and memory*: memory for skills and habits (procedural memory); memory for facts and events (episodic memory); working memory.
- *Space, time and numbers*: spatial navigation and spatial memory; estimation and storage of duration, size and numbers of objects.
- *Multimodal sensory motor integration*: multimodal integration for vision, audition, body representations and motor output.
- *Capabilities characteristic of the human brain*: processing nested structures in language and in other domains (music, mathematics, action); generating and manipulating symbols; creating and processing representations of the self in relation to others.
- *Architectures supporting conscious processing*: brain networks enabling the extraction of relevant information, its broadcast and its maintenance across time; representation of self-related information, including body states, decision confidence, and auto-biographical knowledge.

In each case, the HBP should develop highly structured, easily reproducible experimental paradigms, initially for human adults, but transferrable to infants, and eventually simplified in the form of “localiser” tasks. Each cognitive function should be decomposed and characterised, using high-spatial resolution activity maps acquired with high-field MRI, neural dynamics activity reconstructed from M/EEG data, as well as intracranial electro-corticogram (ECOG) and single-cell recordings in epilepsy patients.

In addition, we propose to recruit a small group of subjects for repetitive scanning (10-20 scanning sessions over three months, repeated every year), so that all of the above functions can be characterised with the same brains, and their geometrical inter-relations understood at high resolution (initially at 3 Tesla, then 7 T and ultimately 11.7 T). The neural signatures of these functions would be correlated with subjects’ anatomy and connectivity. The data generated would be deposited in the INCF Brain Atlas and Brainpedia.

Theory

Objective

It is impossible to conceive a single theory or theoretical approach capable of addressing the full range of scientific challenges addressed by the HBP. Nonetheless, theoretical insights from the creative application of mathematics can make a valuable contribution to many different areas of research, from modelling of low-level biological processes, to the analysis of large-scale patterns of activity in the brain and the formalisation of new paradigms of computation (see Figure 20). Very often these very different problems require similar skills and “mind sets”. The HBP should therefore create a team of theoreticians and mathematicians, with an interest in specific issues raised by the project, a willingness to work with lab scientists and engineers on specific challenges, and an interest in exchanging ideas and methods, while maintaining a pluralistic approach on controversial issues.

The goal of their work should be to provide a solid mathematical and theoretical foundation for the project’s work, creating bridges between descriptions of the brain at different spatial and temporal scales, and developing mathematical descriptions of learning and brain dynamics providing insights into the relations between different levels of brain organisation and their contribution to brain function.

State of the art

Understood as mathematical modelling, theoretical neuroscience has a history of at least a hundred years. In general, theoreticians have focused on models addressing specific levels of brain organisation, for instance, the relation of Hebbian learning to cortical development [45], the recall of associative memories [46], the link of temporal codes and Spike Timing-Dependent Plasticity [47], the dynamics of neuronal networks with balanced excitation and inhibition [48, 49]. In most cases, the output consisted of “toy models” amenable to mathematical analysis and to simulation on small personal computers. What is not clear is how to connect the insights from these models, or how to ground these insights in detailed biophysical observations.

These are key themes in the work of the theoretical neuroscientists who have contributed to the HBP-PS. For example W. Gerstner has shown how to extract parameters for simple neuron models directly from experimental data, or from detailed biophysical models [50, 51]. M. Tsodyks, W. Gerstner, N. Brunel, A. Destexhe, and W. Senn have produced models of synaptic plasticity suitable for integration in models of large-scale neuronal circuitry [52, 53, 54, 55]; W. Gerstner, D. Wierstra, and W. Maass have explored models in which plasticity is modulated by a reward signal [56, 17, 18], a basic requirement for so-called reinforcement learning. N. Brunel has produced models of population dynamics using networks of randomly connected simple neurons [49], an approach exploited by G. Deco to construct models of decision-making [15]. A. Destexhe [57, 58] has investigated the integrative properties of neurons and networks, while W. Maass has studied their underlying computational principles [59, 60].

Theoretical Neuroscience

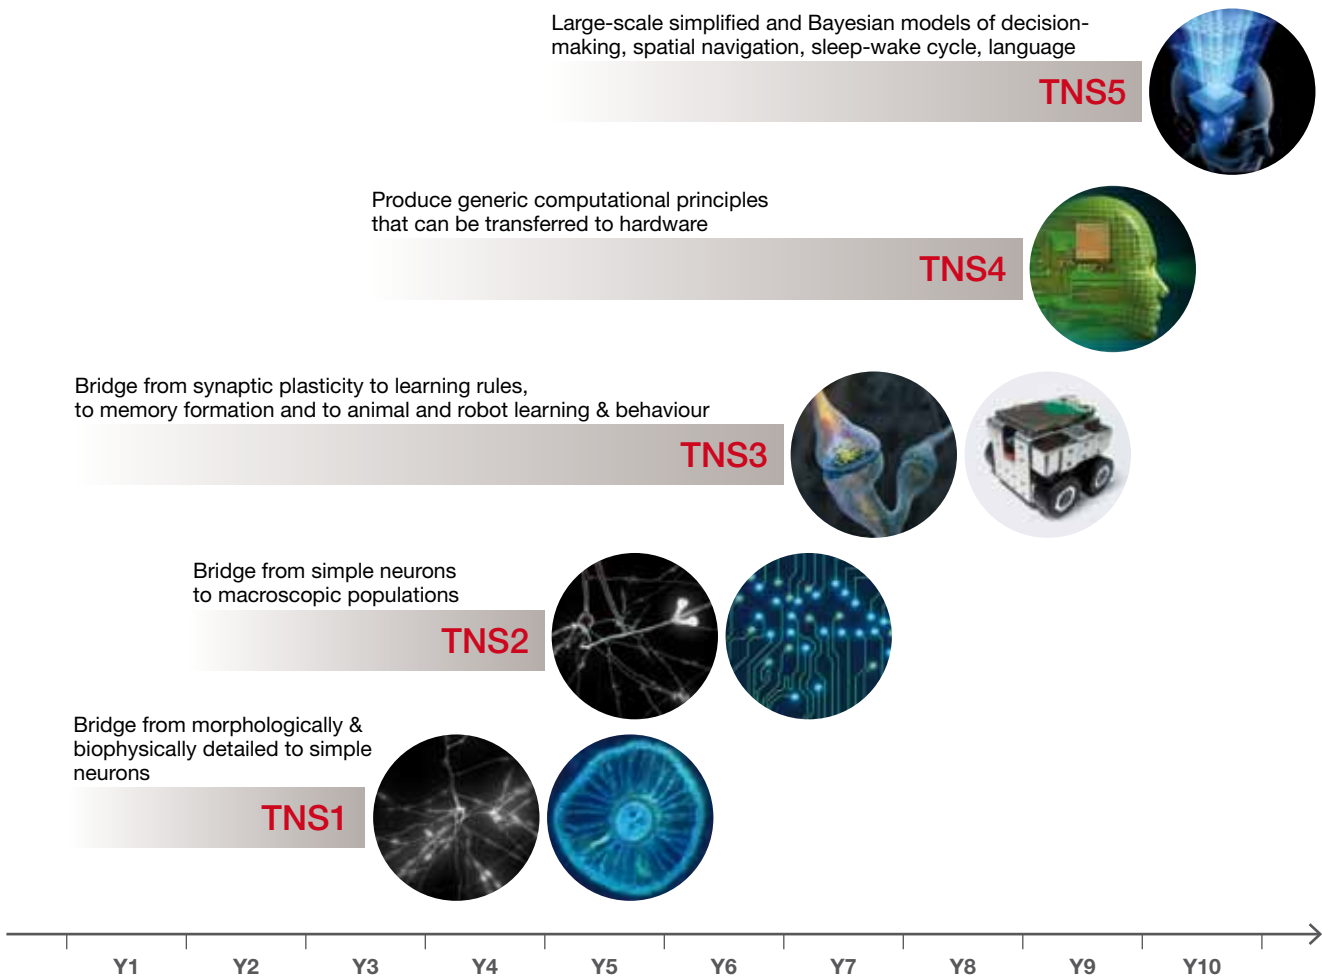


Figure 20: A roadmap for Theoretical Neuroscience (TNS) - major landmarks

Methodology

Theoretical work in the HBP should address a broad range of issues, all related to the goal of achieving a multi-level understanding of the brain.

1. *Bridging scales.* Studies should attempt to establish mathematical principles making it possible to derive simplified models of neurons and neuronal circuits from more detailed biophysical and morphological models, population models and mean field models from simplified neuron models, and brain region models from models of interacting neuronal populations. Other studies should model brain signals at various scales from intracellular signals to local field potentials, VSD, EEG and MEG. The results from this work would provide basic insights into the relationships between different levels of brain organisation, helping to choose parameter values for large-scale modelling, and guiding the simplification of models for implementation in neuromorphic technology.
2. *Synaptic plasticity, learning and memory.* This work should develop learning rules for unsupervised and goal-oriented learning. Key themes would include the derivation of learning rules from biophysical synapse models,

the identification of rules for unsupervised learning and emergent connectivity, rules describing the role of neuromodulation in learning (the role of reward, surprise and novelty), and the functional and medical consequences of disturbances in plasticity on different time scales.

3. *Large-scale brain models.* The HBP should develop simplified large-scale models of specific cognitive functions. These models should provide a bridge between “high-level” behavioural and imaging data and detailed multi-level models of brain physiology. Topics for modelling would include perception-action, multi-sensory perception, working memory, spatial navigation, reward systems, decision-making and the sleep/wakefulness cycle. These models would make a direct contribution to the design of cognitive architectures for neuromorphic computing systems.
4. *Principles of brain computation.* Studies in this area should develop mathematical descriptions of neural computation at the single neuron, neural microcircuit and higher levels of brain organisation. The results would provide basic insights into the multi-level organisation of the brain, while simultaneously contributing to the high-level design of neuromorphic systems.

To encourage collaboration among theoreticians engaged in different areas of theoretical neuroscience, we propose that the HBP creates a European Institute for Theoretical Neuroscience based in Paris. The Institute would run active young researcher, young investigator and visiting scientists programmes and would serve as an attractive meeting point for workshops on topics related to the goals of the HBP.

ICT platforms

Neuroinformatics Platform

Objective

The HBP should build and operate a *Neuroinformatics Platform*, organising data, knowledge and tools from different areas of neuroscience and medicine. These are goals it shares with the INCF [3] and with other on-going projects in large-scale neuroscience – in particular the *Allen Institute's* Brain Atlas projects (www.brain-map.org). The HBP should work with these organisations to develop neuroinformatics tools that facilitate this task. These should include tools for the analysis and interpretation of large volumes of structural and functional data as well as generic tools for the construction of multi-level brain atlases. The HBP should then use these tools as part of a broad international effort to develop detailed multi-level atlases of the mouse and human brains, bringing together data from the literature, and from on-going research, and providing a single source of annotated, high quality data for the HBP modelling effort and for the international neuroscience community (see Figure 21).

Another important goal should be to establish predictive neuroinformatics as a valid source of data for modelling. One of the most important functions of the *Neuroinformatics Platform* should be to identify correlations between data for different levels of biological organisation, making it possible to estimate parameters where experimental data is not available. Systematic application of this strategy has the potential to drastically increase the amount of information that can be extracted from experimental data, rapidly filling gaps in our current knowledge and accelerating the generation of data required for brain modelling.

State of the art

Virtually all areas of modern science face the challenge of providing uniform access to large volumes of diverse data. In neuroscience, with its broad range of experimental techniques, and many different kinds of data, the challenge is particularly severe. Nearly a hundred years of neuroscience research has produced a vast amount of knowledge and data, spread across thousands of journals. The challenge now is to provide uniform access to this data.

The first attempts to achieve this goal date back to 1989, when the Institute of Medicine at the US National Academy of Sciences received funding to examine how information technology could create the tools needed to handle the growing volume and diversity of neuroscientific data. The study

At least thirty million research papers relevant for unifying models of the brain

report, published in 1991 [20] enabled NIMH, to create its own *Human Brain Project*, an effort that lasted until 2004. The work produced many important neuroscience databases. However, it never created a standard interface for accessing the data and provided no specific tools for relating and integrating the data. The creation of interoperable databases using standard descriptions and ontologies remained a goal for future work.

Soon after the NIMH project ended, the Global Science Forum of the Organisation for Economic Co-operation and Development (OECD) initiated the INCF [61]. Since 2005, the INCF has driven international efforts to develop neuroscience ontologies, brain atlases, model descriptions and data sharing, and has played an important role in coordinating international neuroscience research and setting up standards. Other important initiatives in neuroinformatics include the US-based *Neuroscience Information Framework* (NIF) [22], and the *Biomedical Informatics Research Network* (BIRN) [62]. These initiatives are collaborating closely with INCF on issues related to infrastructure, the development of brain atlases, ontology-development and data sharing. Another important initiative was the foundation of the *Allen Institute*, which, since its foundation in 2003, has become a world leader in industrial-scale data acquisition for neuroscience. The *Allen Institute* has developed mouse and human atlases for connectivity, development, sleep and the spinal cord, recently investing an additional \$300M for *in vivo* data acquisition and modelling related to the mouse visual system [24]. This work would contribute directly to the HBP.

A second key area of activity for the HBP would be predictive neuroinformatics, a relatively new area of research. Examples of work in this area include a recently published algorithm that can synthesise a broad range of dendritic morphologies [25], algorithms to generate specific motifs in network connectivity [63], and algorithms to predict synaptic strength based on network architecture [64]. In another area of research, recent work has demonstrated that biophysical models of neurons' electrophysiological properties can successfully predict ion channel distributions and densities on the cell surface [65]. By combining these predictions with cellular composition data, it is possible to predict protein maps for neural tissue. Finally, predictive neuroinformatics can help to resolve one of the most important challenges for modern neuroscience, namely the classification and categorisation of different types of cortical interneurons [66]. A recent model [67] uses gene expression data to predict type, morphology and layer of origin with over 80% accuracy. The same model reveals rules for the combinatorial expression of ion channel genes [68].

Neuroinformatics Platform

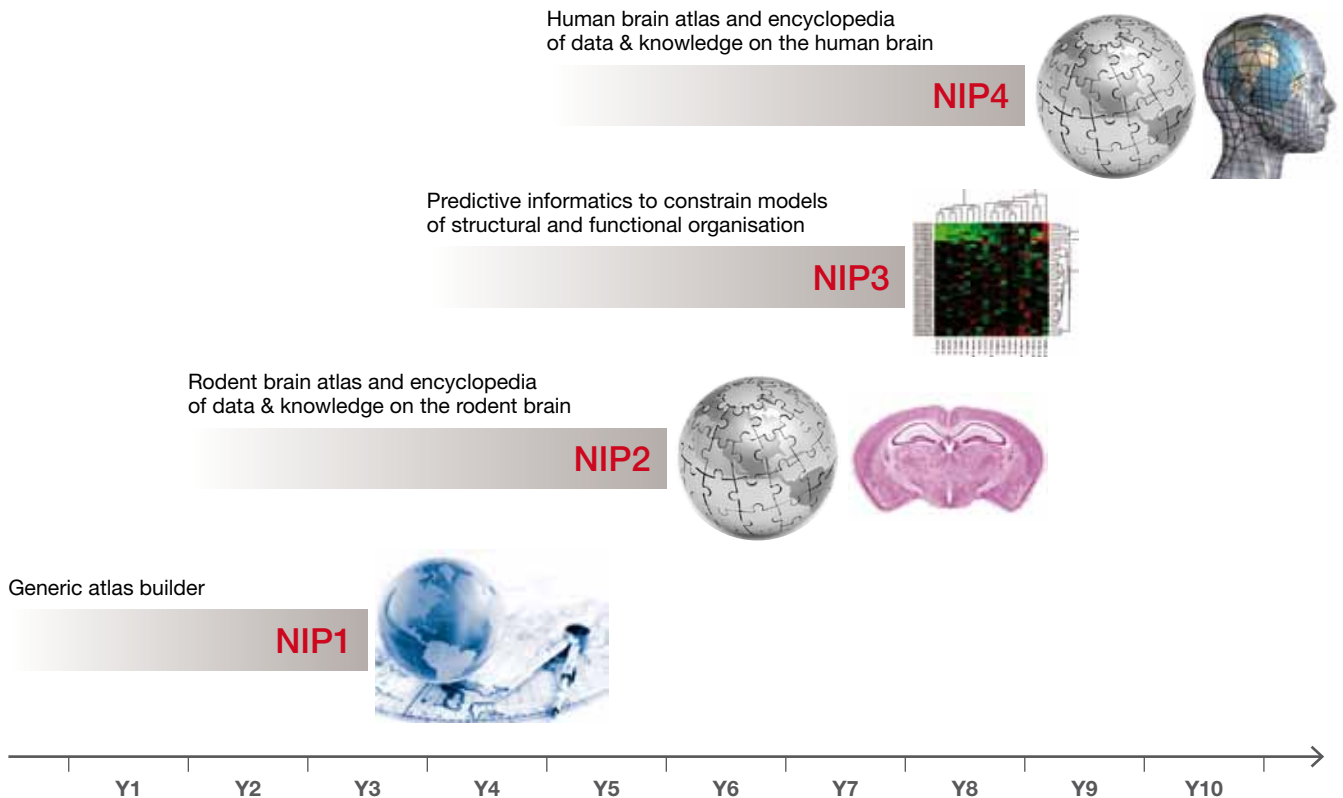


Figure 21: A roadmap for the Neuroinformatics Platform (NIP) - major landmarks

Methodology

The HBP effort in neuroinformatics should focus on five main themes of research.

Tools for brain atlases. The HBP should create a general-purpose software framework, allowing researchers to build and navigate multi-level atlases of the brain of any species. These tools, which would be placed in the public domain, would allow researchers to upload and access information about any part of the brain (identified by its 3D coordinates) at a required level of description. The information contained in the atlases would be distributed across databases in different physical locations. The framework would provide a shared data space, ontologies, data mining tools, standards and a generic “Atlas Builder”, making it possible to build, manage and query such atlases. In addition to this work, the project would also create a “Brainpedia” – a community driven Wiki that provides an encyclopedic view of the latest data, models and literature for all levels of brain organisation.

Tools to analyse data on brain structure. Much of the structural data produced by modern neuroscience takes the form of image stacks from light and electron microscopy, MRI, PET etc. Given that many of these techniques produce terabytes of data in a single session, the best way to unlock the information they contain is through automatic image processing. The HBP should invest a significant ef-

fort in developing the necessary tools, which it would share with the community, via the INCF. The results would include software to automate the extraction of key features and statistics needed to construct models, including cell densities and distributions, reconstruction of neuron morphologies, subcellular properties such as synapse and organelle geometry, size and location, brain morphology, and long range fibre tracts underlying connectivity. Additional tools should make it possible to identify, annotate and integrate data from experiments revealing subcellular, cellular and supracellular structures.

Tools to analyse data on brain function. Understanding of brain function depends on data from a wide range of techniques. It is important that simulation results should be comparable against this data. To meet this need, the HBP should develop new tools and techniques that can be used to compare data from simulations against data from experiments (single neuron recordings, measurement of local field potentials, EEG, fMRI, MEG etc.).

Brain atlases. The tools just described should become part of the HBP’s contribution to INCF atlases of the mouse and human brains. The design should encourage research groups outside the project to deposit data in the atlases. This would enable global collaboration on integrating data across scales about the brain in a unified location for each species.

Predictive neuroinformatics. The HBP should make a major effort to develop new tools for predictive informatics, using machine learning and statistical modelling techniques to extract rules describing the relationships between data sets for different levels of brain organisation.

Brain Simulation Platform

Objective

The *Brain Simulation Platform* should consist of a suite of software tools and workflows that allow researchers to build biologically detailed multi-level models of the brain that integrate large volumes of data spanning multiple levels of biological organisation (see Figure 22). These models would generate “emergent” structures and behaviours that cannot be predicted from smaller data sets. The platform should make it possible to build models at different levels of description (abstract computational models, point neuron models, detailed cellular level models of neuronal circuitry, molecular level models of small areas of the brain, multi-scale models that switch dynamically between different levels of description), allowing experimentalists and theoreticians to build the models most appropriate to the questions they are asking and making it possible to build simplified models of features or areas, where there is not enough data to build a more detailed model. The *Brain Simulation Platform* should be designed to support continuous refinement and automated validation as more data becomes available. In this way, models would become steadily more accurate and detailed as the project proceeds.

The tools made available through the platform should allow researchers to perform *in silico* experiments including systematic measurements and manipulations impossible in the lab. These experiments would provide the basic tools needed to explore the multi-level organisation of the brain, and the way it breaks down in disease. Modelling and simulation should contribute to identifying the neuronal architectures underlying specific brain functions, facilitating the simplification of neuronal circuitry for implementation in neuromorphic technology (see below). The project should use these tools to develop and validate first draft models of different levels of brain organisation, in mice and in humans. These models should lead towards models of whole mouse and human brains, mixing simple and detailed neuron models.

State of the art

Early models of the brain attempted to explain brain functions, such as learning and memory, in terms of the behaviour of neurons and neuron populations, giving rise to the fields of Artificial Neural Networks and machine learning [69]. In parallel, other researchers developed mechanistic models that explained brain functions in terms of biological processes. In particular, Hodgkin and Huxley’s seminal model of the generation of neuronal Action Potentials [70] and Rall’s application of cable theory to signal propagation in dendrites [71] made it possible to build models of the brain from its basic components. Other models helped to understand the dynamics of large networks of excitatory and in-

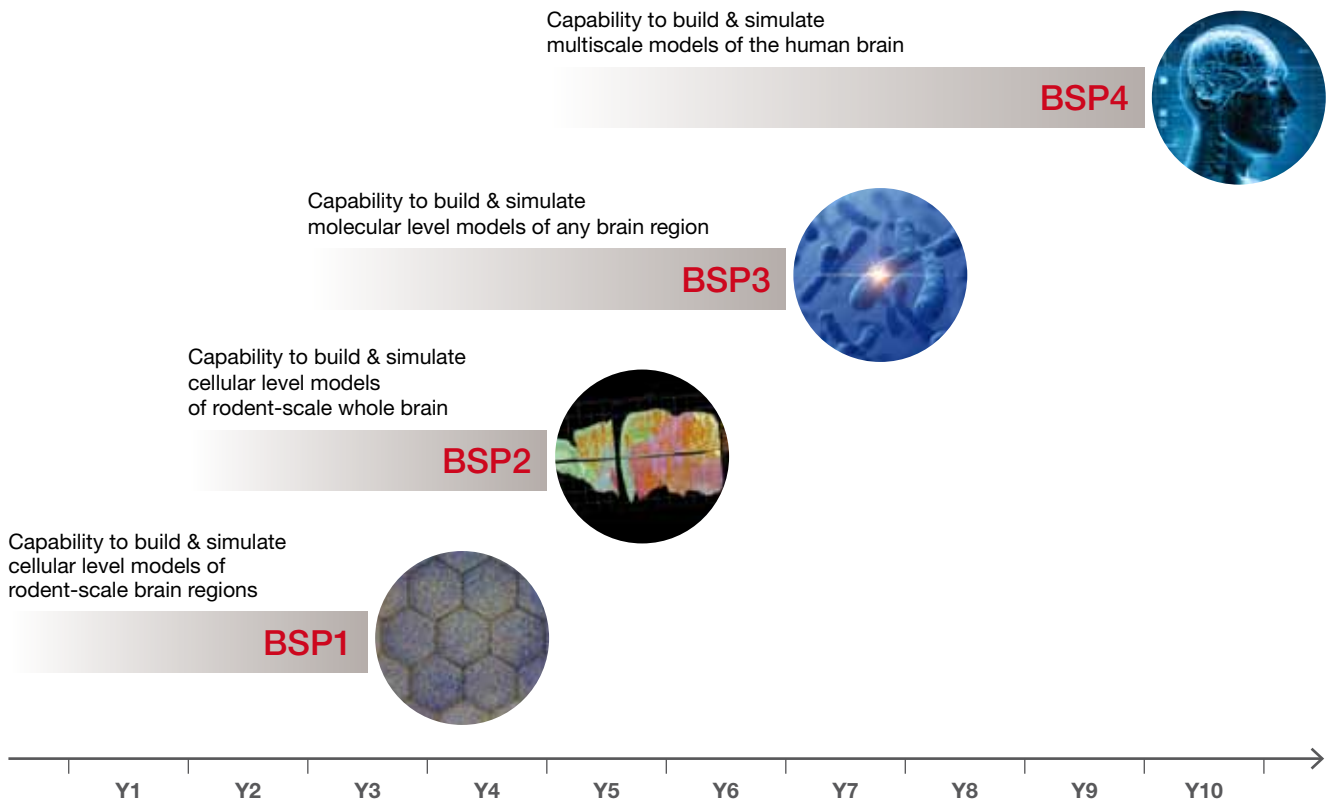
hibitory neurons. In the 1980s, Roger Traub [72, 73] used an IBM 3090 mainframe computer to simulate 10,000 neurons, each with about 20 compartments. Ten years later, De Schutter and Bower pushed the complexity of multi-compartment neuron models to simulate a cerebellar Purkinje cell [74, 75], with over 1600 compartments, and Obermayer et al. pioneered the use of parallel computers for large-scale simulations of simplified neurons [76]. Since then, rapid improvements in supercomputer performance have made it possible to simulate ever-larger models. In 2005, for instance, Izhikevich reported a feasibility study simulating a network with 10^{11} neurons and 10^{15} synapses, numbers comparable to the numbers of neurons and synapses in the human brain. However, each neuron was represented by a single compartment, synapses were not explicitly represented and connections had to be recomputed on each simulation step [77]. In 2007, Djurfeldt et al. reported a large-scale simulation of a columnar cortex with 10^7 detailed multi-compartment neurons and 10^{10} synaptic connections [78]. In the same year, Morrison reported the simulation of a network with 10^9 synapses and spike-timing dependent plasticity (STDP) [79]. In 2009, the Modha group at the IBM Almaden Research Center reported the simulation of a network, with roughly the same numbers of neurons and synapses as the brain of a cat (10^9 neurons and 10^{13} synapses) [80, 81].

In parallel with work on very large-scale networks, many groups have developed general-purpose simulators allowing simulation of the brain at different levels of biological detail. Examples of simulators for large networks of relatively simple neurons include Topografica [82], PCSIM [82], MIIND [83], and NEST [84]. NEURON [78] makes it possible to simulate morphologically complex neurons and networks of neurons, and can be integrated with molecular-scale simulations that add biochemical details to its electrical modelling. STEPS [65], MCELL and Brownian Dynamics simulations bridge the gap between NEURON’s compartment electrical model and the molecular-scale processes of diffusion in complex fluid environments and reaction mechanisms such as ligand binding to receptors. Drug binding events and protein-protein interactions are captured using atomistically-accurate but computationally-demanding molecular dynamics simulations. To date, however, there have been relatively few attempts to integrate models and simulations across multiple levels of biological organisation. This is one of the aims of EPFL’s Blue Brain Project [86], whose work provides a starting point for the kind of modelling that would be pursued by the HBP simulation effort. The Blue Brain Project, which started in 2005, is the first attempt to develop a unifying model of the neocortical column of juvenile rat, based on detailed anatomical and electrophysiological data. A key part of the work of the project has been the development of the necessary software and workflows [87, 88], which would be further developed in the Human Brain Project.

Methodology

The HBP should develop a suite of software tools, workflows and services allowing researchers from inside and outside the project to collaboratively build and simulate detailed

Brain Simulation Platform

Figure 22: A roadmap for the *Brain Simulation Platform* (BSP) - major landmarks

models of the brain, at the level of detail best adapted to the questions they seek to answer. These would be “snap-shot” models, representing the multi-level structure of a brain at a given stage in its development. Initial parameter values would be based on statistical data from experiments and predictive neuroinformatics and validated against data from biological experiments. We hypothesise that by “training” such models in closed-loop set-ups (see page 49), it would be possible to build systems displaying realistic behavioural and cognitive capabilities.

To achieve these goals the *Brain Simulation Platform* should provide the following functionality.

1. *Brain Builder: A software engine for detailed brain models.* The Brain Builder should make it possible to build brain models of any species, at any age at any desired level of detail, so long as the necessary data is available. The same software should make it possible to build models incorporating hypotheses of disease causation (e.g. absence of specific ion channels or receptors, pathological patterns of network connectivity). The Brain Builder should include tools to embed data from brain atlases (see above), a “multi-scale slider” allowing modellers to vary the resolution of their models, tools to set up closed-loop experiments and tools to deploy simulations to high performance computing platforms.
2. *Brain simulation engines.* The platform should incorporate a multi-scale simulation framework integrating existing simulation engines for molecular, cellular, and network level simulation. The framework should make it possible to build and simulate models representing the same tissue at different scales, laying the foundations for studies of the relations between levels.
3. *Molecular dynamics simulations.* The HBP should use molecular dynamics simulations and a range of coarse-grained techniques to generate molecular level information for the project’s multi-scale models. The project would use this information to improve models of ion channels and receptors, to create coarse-grained models of the dynamics of cell membranes and organelles, to understand protein-protein interactions and to understand the way drugs bind to proteins. The same information would provide vital input for the development of coarse-graining strategies for large-scale molecular simulations.
4. *Brain models.* The platform should incorporate first draft models representing different levels of brain organisation (molecular level models of selected neurons, neuromodulation and synapses, synaptic plasticity and homeostasis, glia and neuro-vascular coupling, cellular level models of major classes of neurons, and of neural microcircuits in important regions of the brain, cellular level models of whole brain regions and brain systems, mixing point neurons and detailed neuron models). These should lead to models of whole mouse and human brains, which would exploit the platform’s multi-scale capabilities.

5. *The Brain Simulation Platform*. The platform should be accessible to researchers via an Internet-accessible Brain Simulation Cockpit, providing them with the tools they need to build brain models, set up *in silico* experiments, and analyse the results. The platform should allow them to perform *in silico* experiments investigating the relationships between different levels of biological organisation in the healthy and the diseased brain and preparing the way for the re-implementation of neuronal circuits in neuromorphic hardware (see below). A professionally managed service centre should provide them with the necessary support and training. The project should support a vigorous visitors programme for external scientists wishing to make use of these services in their own research.

High Performance Computing Platform

Objective

Current supercomputing technology lacks the computing and communication power, I/O capabilities and advanced software necessary for multi-scale modelling of the human brain. Just as importantly, it also lacks the software capabilities to analyse and visualise the massive volumes of data that will be produced by large-scale brain simulations. The goal of the *High Performance Computing Platform* should be to fill these gaps, providing the project and the wider community with the advanced supercomputing capabilities they require. The platform should consist of a main *Human Brain Supercomputing Facility* that gradually evolves toward the exascale over the duration of the project (see Figure 23). This should be complemented by satellite facilities dedicated to software development, molecular dynamics simulations, and massive data analytics. A key goal should be to develop a capability for *in situ* analysis and visualisation of exascale data sets and for interactive visual “steering” of simulations. This kind of *interactive supercomputing* would be invaluable not just for brain simulation but also for a broad range of other applications, in the life sciences and elsewhere.

State of the art

Much, though not all high performance computing relies on so-called supercomputers – computers that perform at or near the highest speed possible in a given period, usually measured in *Floating Point Operations per Second* (“flops”). Today’s supercomputers are generally massively parallel systems, sometimes containing hundreds of thousands of interconnected processor cores. The largest machines achieve peak performances of several Pflops (10^{15} flops) [89].

Since the introduction of the first supercomputers by Cray in the 1960/70s, trends in supercomputer performance and memory have followed “Moore’s Law”, according to which the performance of computer chips doubles approximately every eighteen months. The International Technology Roadmap for Semiconductors (ITRS) [90] foresees that this scaling will continue for several chip generations to come. However, even this very rapid improvement in performance will not be sufficient to achieve exascale computing this de-

cade. It appears that the only way to reach this goal - already announced by major manufacturers - is by further increasing the number of processors in each machine.

This strategy poses severe technical challenges [91, 92]. For environmental and business reasons, vendors have set themselves the goal of containing energy consumption to a maximum of 20 megawatts. This requirement is driving processor design in the direction of power-efficient many-core CPUs, playing a role similar to today’s GPUs. On the system and application software side, the massive parallelism of exascale machines will raise major issues of programmability and resilience to errors. Memory and I/O constraints will present additional obstacles. Given the current state of the art, it is unlikely that memory capacity and communications bandwidth will keep up with the expected increase in compute performance. In these conditions, energy considerations will make it prohibitively expensive to move large amounts of data from system memory to hard disk, the current practice for offline analysis of simulation results.

International supercomputer vendors like IBM and Cray and exascale research initiatives are making intensive efforts to solve these problems [93, 94]. In October 2011, the European Union announced the funding of CRESTA [95], DEEP [96] and Mont-Blanc [97], three complementary projects each studying different aspects of the exascale challenge. The HBP should collaborate with these projects, ensuring that the technology developed meets the requirements of brain simulation, which in some respects are qualitatively different from those of other applications, for example in physics.

Ever since the pioneering work of Gerstein and Mandelbrot in the 1960s [98], brain simulation has used the latest computing hardware available. This tendency continues today as teams in the USA, Europe, and Japan work to increase the power of simulation technology. In the USA, many of these efforts are coordinated by the DARPA SyNAPSE programme [99]. In Japan, efforts to simulate the whole brain are funded by the MEXT “Next Generation Supercomputer” project [100]. In Europe, the EU-funded BrainScaleS [6] and the UK-funded SpiNNaker [101] projects are working to enable multi-scale simulations of the brain on custom neuromorphic hardware.

These projects mainly focus on models with large numbers of neurons and synapses but with little or no detail at lower levels of biological organisation. The HBP, by contrast, would build and simulate biologically realistic models of the complete human brain, at least at the cellular level, and use them as the basis for *in silico* experiments. EPFL’s on-going Blue Brain Project (BBP) [102], which has pioneered this approach, has produced a parallel version of the NEURON code, running on an IBM Blue Gene/P supercomputer with a peak performance of 56 Tflops. This is sufficient to run cellular-level models with up to 1 million detailed, multi-compartment neurons. A simple extrapolation suggests that after optimisation, a large Blue Gene/P system such as the 1 PFlop machine at the Jülich Supercomputing Centre could simulate up to 100 million neurons – roughly the number found in the mouse brain. Cellular-level simulation of the

High Performance Computing Platform

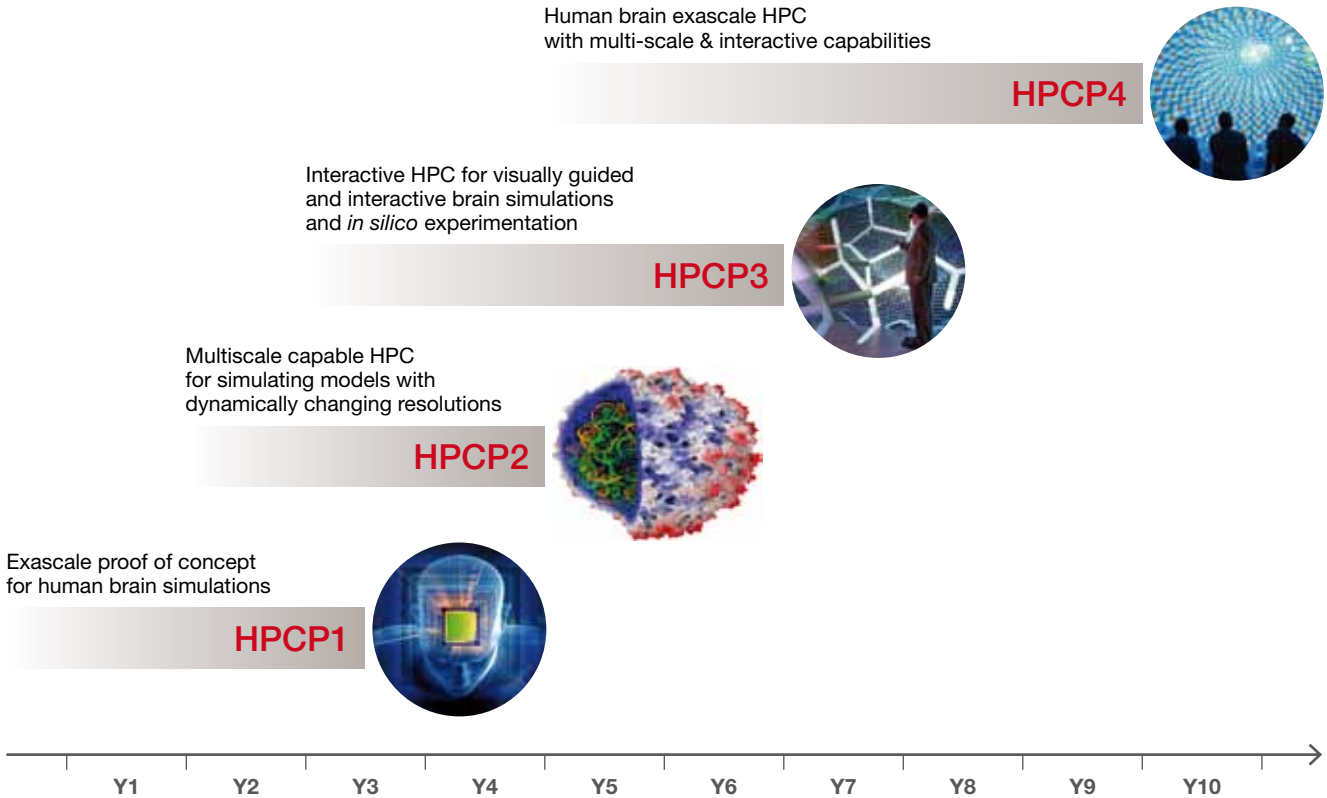


Figure 23: A roadmap for the High Performance Computing Platform (HPCP) - major landmarks

100 billion neurons of the human brain will require compute power at the exascale (10^{18} flops).

A second unique requirement of the Human Brain Project is that supercomputing hardware should act as an interactive scientific instrument, providing researchers with visual feedback and allowing them to “steer” simulations while they are underway. This is very different from the *batch mode* in which most supercomputers are operated today. Creating this capability will require completely new developments in supercomputing software, including new techniques for *in situ* visualisation and data analysis.

Methodology

Building and simulating multi-level models of the complete human brain will require exascale supercomputing infrastructure with unprecedented capabilities for interactive computing and visualisation. We recommend that the HBP should work with European exascale research projects and leading manufacturers to develop the necessary software and hardware.

1. *Developing exascale supercomputing for brain research.* The HBP should collaborate with major international manufacturers (IBM, Cray) and with exascale research initiatives, like DEEP and Mont-Blanc, that include European HPC manufacturers (EuroTech, Bull). The end goal should be to design and deploy the supercomputing technology required by the project, gradually moving towards exascale capabilities, expected to be available by 2020.

2. *Numerical methods, programming models and tools.* To support efficient interactive simulation of brain models, the project should develop new numerical methods, parallel programming models, and performance analysis tools adapted to the extreme parallelism of future exascale systems and should also develop new middleware for workflow and I/O management.
3. *Interactive visualisation, analysis and control.* The project should develop a novel software framework allowing interactive steering and *in situ* visualisation of simulations. The development work should produce both general-purpose software and neuroscience-specific interfaces – *virtual instruments* allowing scientists to work with virtual specimens in the same way they work with biological specimens.
4. *Exascale data management.* HBP brain simulations would generate massive amounts of data. An important task should thus be the design and development of technology making it possible to manage, query, analyse and process this data, and to ensure that it is properly preserved.
5. *The High Performance Computing Platform.* The HBP High Performance Computing Platform should make the project’s supercomputing capabilities available to the project and the community. The platform should consist of a production-scale *Human Brain Supercomputing Facility* at Jülich, a smaller software development system at CSCS, Switzerland, a system for molecular-level

simulations at Barcelona Supercomputing Center, and a system for massive data analytics at CINECA, Italy. The four systems should be connected via a dedicated fast network. Data storage should be provided directly by the centres and through cloud services. The project should also provide user support and training, coordination with PRACE and other research infrastructures, cooperation with industry, and a scientific visitors programme.

Medical Informatics Platform

Objective

The goal of the *Medical Informatics Platform* should be to provide the technical capabilities to federate imaging and other clinical data currently locked in hospital and research archives and databases while guaranteeing strong protection for sensitive patient information. These capabilities should include tools to search for, query and analyse the data. The platform should make tools and data available to the clinical research community, using them to develop a comprehensive classification of brain diseases, based on parameterised combinations of biological features and markers. Success in this enterprise would accelerate the development of a new category of biologically based diagnostics, supported by strong hypotheses of disease causation. The *Brain Simulation Platform* would enable *in silico* experiments to test these hypotheses. If successful, this effort would lead to the identification of new drug targets, and other strategies for the treatment of brain disease (see Figure 24). Brain simulation should make it possible to predict their desirable and adverse effects, providing valuable input for industry decision-makers before they invest in expensive programmes of animal experimentation or human trials.

State of the art

Recent years have seen very little progress in the development of new drugs for brain diseases. This is partly due to difficulties in diagnosing patients. Alzheimer's disease, for instance, is misdiagnosed in as many as 20% of cases [103, 104]. However, the main reason is the lack of detailed causal explanations of the way diseases come about and the factors determining their manifestations. It is essential, therefore that researchers take full advantage of advances in genetics, imaging, database management and supercomputing, integrating and exploiting data of different types and from different sources.

Data sharing among clinical scientists is less common than in other scientific communities. According to Visscher et al. [105], the reasons include the need for standardisation, the time required to transfer data to repositories, the need to protect clinical confidentiality, the perceived risk of jeopardising publications, and difficulties in assessing the accuracy of results. All these problems are soluble in principle, and have already been solved by other scientific communities.

Imaging presents an illustration of the challenges and potential solutions. European hospitals and research estab-

Biological signatures of disease will drive a paradigm shift from symptom to biologically based classifications of disease

lishments generate an enormous number of brain images for clinical purposes, most of which are only viewed once before being archived on hospital or laboratory servers. Much of this data consists of structural (sMRI) images scanned at 1.5 or 3.0 Tesla. The variance introduced by averaging image data from multiple imaging platforms is less than the variance attributable to disease [106]. This suggests that archived images represent a largely unused resource for population-based investigations of brain diseases.

Several attempts to exploit such data are already in progress. Thus, grant-awarding institutions such as the NIH and Wellcome Trust require databases to be made public on the Internet, facilitating data sharing. Switzerland, among other countries already allows hospital data mining by health economists and insurance companies to improve the quality of health care. Pilot studies by partners in the HBP-PS are profiting from this favourable situation to mine anonymised patient data collected by pharmaceutical firms, including data from failed clinical trials.

Preliminary international data generation initiatives, such as the ADNI database [107] have demonstrated practicability and value for money, informing a broad range of experiments conceived, executed and published independently by internal and external collaborators.

Methodology

The *Medical Informatics Platform* should build on existing international data generation initiatives, allowing researchers to query and analyse large volumes of clinical and other data stored on hospital and laboratory servers. The work required to build and manage the platform can be summarised as follows.

Federated data management. The *Medical Informatics Platform* should provide a software framework allowing researchers to query clinical data stored on hospital and laboratory servers, without moving the data from the servers where it resides and without compromising patient privacy. The data made available through the platform should include brain scans of various types, data from electrophysiology, electroencephalography and genotyping, metabolic, biochemical and haematological profiles, data from validated clinical instruments used to quantify behaviour and emotion as well as relevant data on provenance.

Data acquisition and integration. The HBP should recruit hospitals, research labs, industrial companies and other large-scale data gathering initiatives (e.g. large longitudinal

Medical Informatics Platform

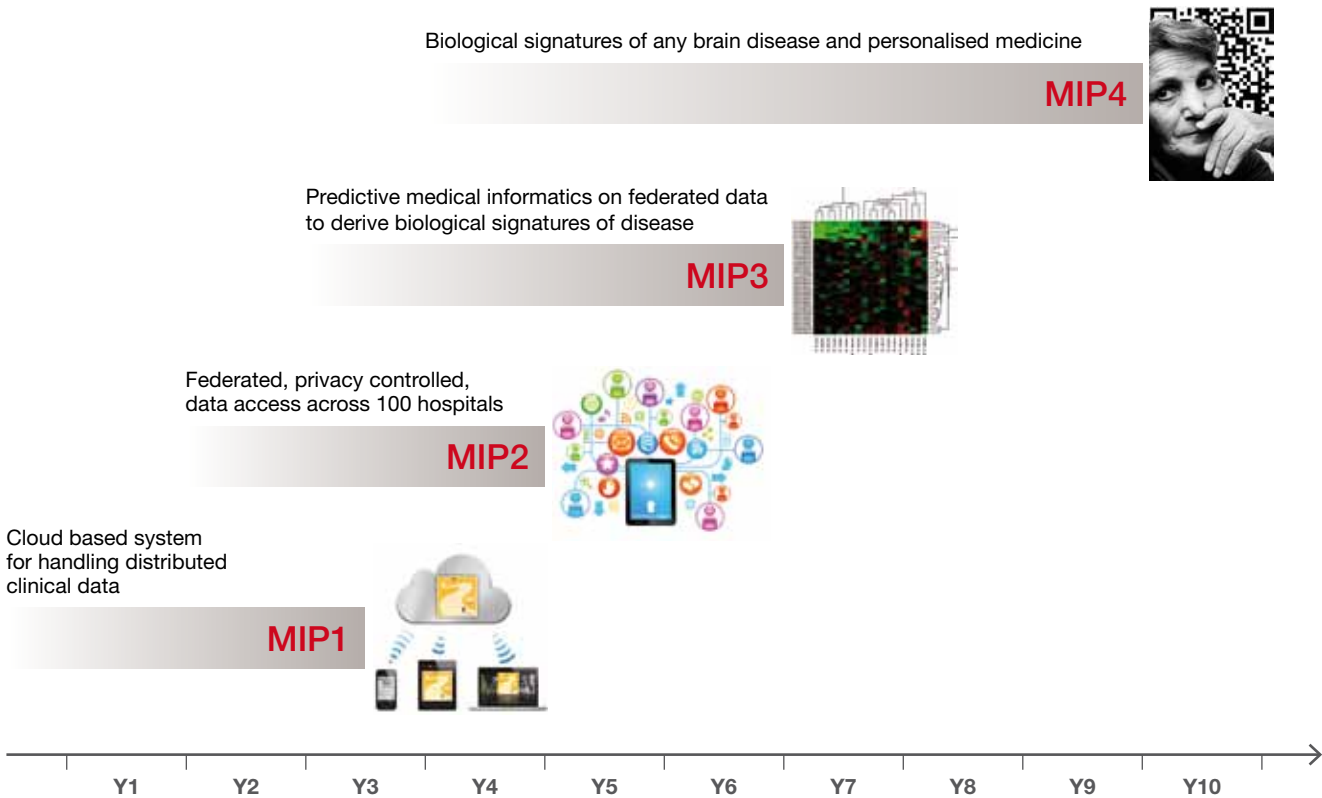


Figure 24: A roadmap for the *Medical Informatics Platform (MIP)* - major landmarks

studies) to make their data available through the platform. As part of this work it should develop common protocols for data capture. A key goal should be to move away from a culture of data protection to one of data sharing, exploiting the robust data protection offered by the platform. To create incentives for this change, the HBP should ensure that researchers who contribute data to the platform have free access to data contributed by other researchers. Such a policy would encourage greater efficiency in the use of data, stronger collaboration among researchers, the creation of larger cohorts, and more effective approaches to rare diseases.

Medical intelligence tools. The HBP should build software tools and techniques of data analysis, making it easier for researchers to analyse data made available through the platform. These would include machine learning tools, Bayesian model selection techniques and high dimensional data mining algorithms. The tools and techniques made available by the project should be used to study clinical and experimental data for the full range of brain disease, detecting recurrent patterns (biological signatures) associated with specific diseases. Focused scientific investigation should make it possible to associate biological signatures of diseases with differential sites of dysfunction, abnormal development or disorganisation, providing evidence of causative mechanisms and allowing the identification of potential drug targets.

Building and operating the platform. The *Medical Informatics Platform* should offer researchers tools to contribute data to the platform and to analyse and exploit the data it provides. Researchers using the platform should receive all necessary training, technical support and documentation.

Neuromorphic Computing Platform

Objective

The HBP should design and implement a *Neuromorphic Computing Platform* that allows non-expert researchers to perform experiments with Neuromorphic Computing Systems (NCS) implementing cellular and circuit-level brain models. The devices would incorporate state-of-the-art electronic component and circuit technologies as well as new knowledge arising from other areas of HBP research (data generation, theory, brain modelling). The platform would provide hardware implementations of large-scale models running in real time (numerical models running on digital multicore architectures), in accelerated mode (physical emulations of brain models) and on hybrid systems (see Figure 25). The platform would be tightly integrated with the *High Performance Computing Platform*.

State of the art

The primary challenges for traditional computing paradigms are energy consumption, software complexity and component reliability. One proposed strategy for addressing these chal-

Neuromorphic Computing Platform

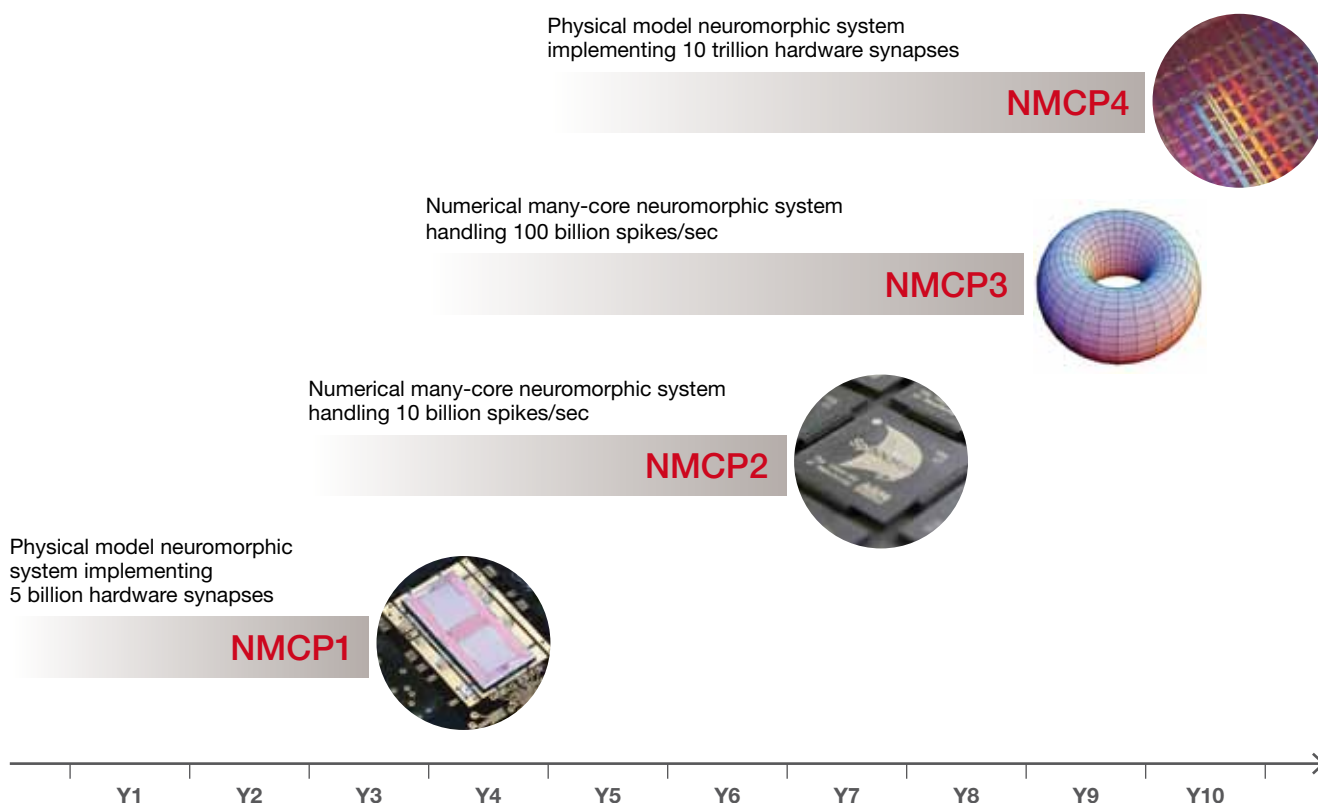


Figure 25: A roadmap for the *Neuromorphic Computing Platform* (NMCP) - major landmarks

lenges is to use *neuromorphic technologies* inspired by the architecture of the brain. Such technologies offer an energy cost per neural operation that is many orders of magnitude lower than the equivalent cost for brain models running on conventional supercomputers (see Figure 15). Other advantages include support for plasticity and learning and the ability to run at speeds up to a thousand times faster than biological real time, making it possible to emulate the dynamics of model systems over periods from milliseconds to years (see Figure 27). Neuromorphic architectures exploit the characteristics of inherently noisy and unreliable nanoscale components with characteristic sizes approaching the atomic structure of matter. This is not possible with classical computing architectures.

Neuromorphic computing was pioneered by the group of Carver Mead [108] at Caltech, the first to integrate biologically inspired electronic sensors with analogue circuits and to introduce an address-event-based asynchronous, continuous time communications protocol. Today, the Mead approach is followed by many groups worldwide, notably the Institute for Neuroinformatics at ETH Zürich (Switzerland) [109].

The main focus of the Mead work is on the demonstration of basic computational principles. By contrast, IBM's *SyNAPSE (Systems of Neuromorphic Adaptive Plastic Scalable Electronics)* project aims to reproduce large systems that abstract away from the biological details of the brain and focus on the brain's larger-scale structure and architecture – the way its elements receive sensory input, connect to each other, adapt these connections, and transmit motor output. The

proponents of this approach argue that its inherent scalability would allow them to build systems that match the computing efficiency, size and power consumption of the brain and its ability to operate without programming [99].

The European FACETS project has pioneered a different approach that combines local analogue computation in neurons and synapses with binary, asynchronous, continuous time spike communication [110-112]. FACETS systems can incorporate 50×10^6 plastic synapses on a single 8-inch silicon wafer. BrainScaleS – a follow-up project – is pioneering the use of the technology in experiments that emulate behaviour and learning in closed-loop experiments over periods of up to a year while simultaneously emulating the millisecond-scale dynamics of the system. In the near future, advances in CMOS feature size, connection technologies and packaging should make it possible to build multi-wafer systems with 10^{13} plastic synapses operating at acceleration factors of 10,000 compared to biological real-time. The FACETS group has also pioneered a unified concept for a network description language (PyNN) that provides platform independent access to software simulators and neuromorphic systems [113].

Yet another strategy for neuromorphic computing is to implement brain models in classical many-core architectures. This is the approach adopted by the SpiNNaker group in the UK [114, 115]. The group has a strong grounding in the ARM architecture, which offers an excellent basis for scalable digital many-core systems operating at real time with low power. The project has recently completed the integration of a

Neuromorphic techniques can produce low-cost, energy-efficient computing systems exploiting the properties of inherently noisy, unreliable component technologies, down to the nano-scale

SpiNNaker chip into an operational system and is now running experiments [116, 117]. Each of these chips has 18 cores and a shared local 128M byte RAM, and allows for real-time simulation of networks implementing complex, non-linear neuron models. A single chip can simulate 16,000 neurons with eight million plastic synapses running in real time with an energy budget of 1W.

Methodology

The HBP's strategy for developing neuromorphic computing systems integrates neuromorphic hardware with brain models at different levels of biological detail. The HBP should systematically study the relationship between the computational performance of the neuromorphic systems, and the complexity of the models (model neurons, model connection networks), identifying strategies to reduce complexity while preserving computational performance. Many of these strategies would rely on high performance computing.

The *Neuromorphic Computing Platform* should bring together two complementary approaches to neuromorphic computing – the first based on non-classical, physical emulation of neural circuits, the second on a classical, programme-based many-core approach.

The platform should allow researchers to access the following capabilities.

1. *Neuromorphic computing through physical emulation of brain models.* The project should implement physical emulations of brain cells, circuits and functions in mixed-signal VLSI hardware that builds on technology developed in the FACETS project. The circuits would run up to 10,000 times faster than real time. This capability would allow experiments requiring systematic exploration of parameter space (e.g. to estimate parameter values) or simulation of learning and development over long periods – months or years – of biological time.
2. *Neuromorphic computing with digital many-core simulation of brain models.* In parallel with this work, the project should develop models implemented on scalable many-core digital ASICs. The devices would offer on-chip floating point operations and memory management, as well as fast lightweight packet switched networks, making it possible to develop real-world applications operating in real time (controllers for robots, systems for applications outside robotics). They are thus a crucial element in the HBP's strategy to transform computing technology.
3. *Common software tools and HPC integration.* The platform should include a suite of software tools to support the design and development of neuromorphic systems and applications. These should include tools to import and simplify detailed brain models, tools to develop *Executable Systems Specifications*, and tools to measure the performance of physical and simulated systems.
4. *Novel technologies for neuromorphic circuits.* The HBP should also investigate new hardware approaches to the implementation of neuromorphic circuits. Candidates include new technologies for distributed memory, nanoscale switches, high-density assembly technologies, 3D silicon integration, and novel design methodologies for Neuromorphic VLSI. Where appropriate, the project should develop functional demonstrators, as a first step towards integrating the technologies in the platform.
5. *The Neuromorphic Computing Platform.* The *Neuromorphic Computing Platform* should integrate these tools and technologies in a fully operational Neuromorphic Computing System and make them available to scientists from outside the HBP. The required work would include development of the hardware and software architecture for the platform, development of components (classical and non-classical systems), assembly of components, services to operate and maintain the platform and services to provide users with training and technical support.

Neurorobotics Platform

Objective

The HBP should develop a *Neurorobotics Platform* allowing researchers to set up closed-loop robotic experiments, in which a brain model is coupled to a simulated body, interacting with a simulated world.

State of the art

In previous work with closed-loop robotic experiments, researchers have studied a wide array of topics. Typical examples include cricket phonotaxis [118], the role of place cells in rodent navigation [119], motion control by the cerebellum [120] and categorisation processes involving large areas of the brain [121]. Nonetheless the use of robots in cognitive neuroscience is still relatively rare, partly due to practical difficulties in implementation.

In most current research, the same group takes responsibility for the robot controller, the robot body, and the environment, implementing them with differing degrees of detail and accuracy depending on expertise and interests. Robot controllers have typically been implemented as artificial neural networks with architectures derived from experimental data (see for example [118, 122]). Often, however, the robot body has received less attention and has included only the basic sensing and actuation capabilities needed to perform a particular experiment. For example, [122] reports an artificial rat whose sensor apparatus is limited to rudimentary visual, odometric and short-range proximity sensing; similarly [123] describes the use of a simple robotic arm with a camera

Neurobotics Platform

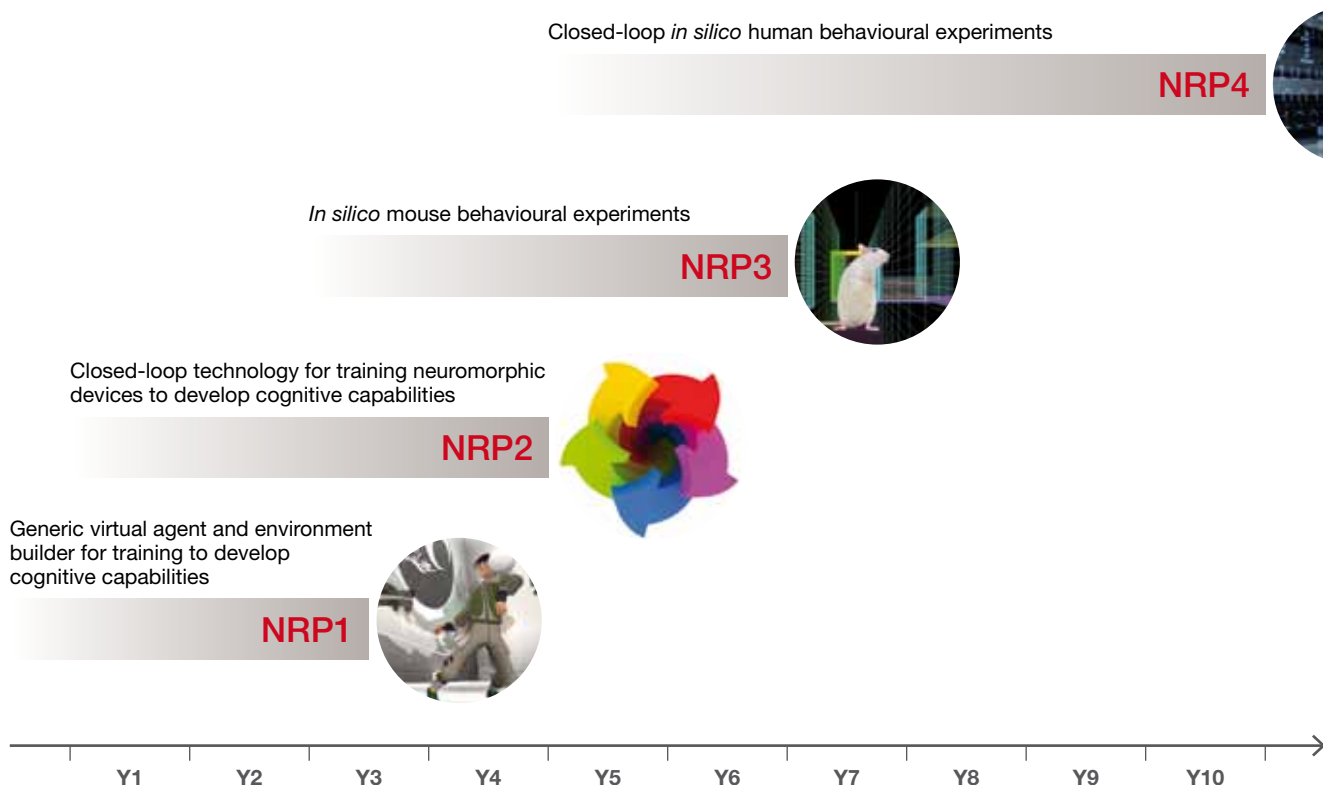


Figure 26: A roadmap for the *Neurobotics Platform* (NRP) - major landmarks

to solve a what and where task. In most cases, the systems developed were used only in-house and were never validated or reused by other scientists.

By contrast, the *Neurobotics Platform* should offer scientists and technology developers a software and hardware infrastructure allowing them to connect pre-validated brain models to detailed simulations of robot bodies and environments and to use the resulting neurobotic systems in *in silico* experiments and technology development work.

Methodology

The HBP *Neurobotics Platform* should make it easy for researchers to design simulated robot bodies, to connect these bodies to brain models, and to embed the bodies in dynamic simulated environments. The resulting set-ups should allow them to perform *in silico* experiments, initially replicating previous experiments in animals and human subjects, but ultimately breaking new ground.

The platform should provide researchers with access to simulated brain models running slower than real time, and to emulated models running faster than real time. In initial experiments, robots and environments would be simulated. For applications-related work requiring real-time operation, the platform should provide access to many-core implementations suitable for use with physical robots and machinery, together with the necessary interfaces (see Figure 26).

The *Neurobotics Platform* should consist of three core modules.

Simulated robots

This module would allow researchers to build simulated robots based on detailed specifications. It would include the following components.

- *A Robot Builder:* a generic tool to design, develop and deploy simulated robots.
- *A Sensory System Builder:* a tool to generate models of perception in different modalities (auditory perception, visual perception etc.).
- *A Motor System Builder:* a tool to generate models of motor systems (muscles or motors) and of the peripheral nervous system.
- *A Brain-Body Integrator:* automated routines for the calibration of brain models to work with the selected body sensory and motor systems.

Simulated environments

This module would allow researchers to build rich simulated environments in which they could test their robots and run experiments. The module would provide the following tools.

- *An Environment Builder:* a generic software tool for designing and deploying dynamically changing simulated environments.
- *An Experiment Designer:* a tool to configure experiments and to specify testing and measuring protocols.

- *An Electronic Coach*: a software tool allowing researchers to define and execute multi-stage training protocols for robots (specification of timing, stimuli, correct and incorrect behaviours, and reward signals for each stage).

Closed-loop engine

This module would make it possible to create a closed loop between a simulated or a physical robot and a brain model. It would include the following.

- *A Closed-Loop Engine*: generic tools to couple software and neuromorphic brain models to simulated and physical robots and to other devices.
- *The Human Interaction Interface*: a software tool that allows human experimenters to interact with robots and their environment.
- *A Performance Monitor*: a set of tools to monitor and analyse the performance of the neurobotic system in its environment, and to produce configurable diagnostic messages.

Building and operating the platform

The HBP *Neurorobotics Platform* would integrate the three modules and provide researchers with a control centre, where they could configure, execute and analyse the results of neurorobotics experiments. A dedicated team would provide users with the training support and documentation required to make effective use of the platform. The HBP should run an active visitors programme for scientists wishing to use the platform.

Applications

Using HBP capabilities to reveal integrative principles of cognition

Objectives

Perhaps the most difficult challenge facing modern neuroscience is the need to account for the causal relationships linking the basic constituents of the brain (genes, molecules, neurons, synapses, microcircuits, brain regions and brain systems) to perception, cognition and behaviour. The HBP’s work in this area should demonstrate that the HBP’s ICT platforms can make a valuable contribution.

To achieve this, the project should fund research projects in which researchers use the platforms to dissect the biological mechanisms underlying specific cognitive and behavioural capabilities and investigating issues of crucial theoretical importance, such as learning and memory, the mechanisms through which the brain represents information (the neural code or codes), and the neural foundations of consciousness and awareness.

Pilot projects should be performed by HBP partners with the necessary know-how and experience. However, the majority of this work should be entrusted to groups and researchers who are not involved in the initial stages of the project, selected via an open competitive process.

Cognition arises from interactions among molecules, cells, and brain regions on time scales spanning 18 orders of magnitude

Time scales

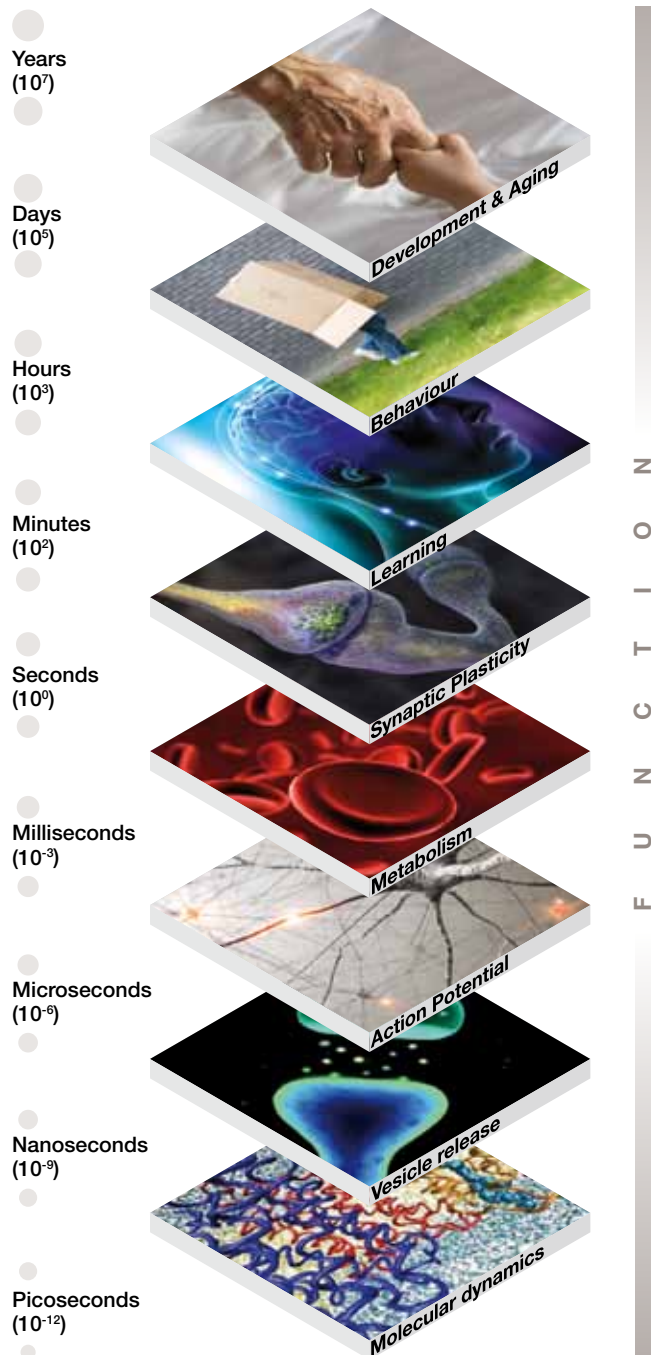


Figure 27: From molecular dynamics to development and aging: time scales span nineteen orders of magnitude

State of the art

The evolutionary function of a brain is to control organisms' behaviour in their environment. In principle, therefore, the only way to test or characterise the high-level behavioural or cognitive capabilities of a brain model is to create a closed loop between the model and a body acting in an environment and to interrogate the model through well-designed experiments (see Figure 28). Once a set-up has successfully replicated we can then identify causal mechanisms by lesioning or manipulating specific brain regions, transmitter systems, types of neuron etc.

Although robotics has yet to win broad recognition as a valid tool for cognitive and behavioural research, a number of groups have attempted to use robots as an experimental tool. Current work can be roughly divided into models based on ideas, models driven exclusively by behavioural data and models that combine biological and behavioural data. It is this last category of model, which is most relevant to the HBP.

An interesting example is Barbara Webb and Henrik Lund's work on cricket phonotaxis [118]. In this pioneering study, the two researchers built an artificial neural network (ANN) reproducing known features of the neuronal circuits believed to be responsible for the female response to the male mating song. Other studies with a similar approach have simulated the role of place cells in rodent navigation [119], motion control by the cerebellum [120] and categorisation processes involving large areas of the brain [121].

Methodology

An initial strategy for closed-loop experiments in the HBP could be defined as follows.

1. Researchers would choose a cognitive or behavioural capability that has already been well characterised in cognitive and behavioural studies and for which theory has already identified a putative cognitive architecture. They would then design an *in silico* experiment to test the ability of a model brain to reproduce this capability and to dissect the multi-level mechanisms responsible. The experimental design would be comparable to the design for an animal or human experiment.
2. They would then use the *Neurorobotics Platform* to design a simulated robot body and a simulated environment, linking the body to a brain model on the *High Performance Computing Platform*, chosen to represent an appropriate level of biological detail, for instance a biologically detailed model for a study of a drug, or a point neuron network for a study of the neuronal circuitry responsible for a particular behaviour.
3. Once the brain model was established, the platform would export a simplified version to a physical emulation of the model running on neuromorphic hardware. The platform would provide the interface to couple the neuromorphic device to the simulated robot and environment. The new set-up would run many times faster

Use Case 1: Tracing causal mechanisms of cognition




Figure 28: Use of the HBP platforms to study mechanisms and principles of cognition: comparing simulated and biological humans and mice on the same cognitive task (touchscreen approach)

than real time, making it possible to *train* the model over long periods of simulated time.

4. Once trained, the model would be tested, comparing the results against animal or human studies. Quantitative and qualitative differences would be analysed, and the results used to refine the brain model, the robot body and the training protocol.
5. Once the model displayed the desired cognitive or behavioural capability, researchers would dissect the underlying neural mechanisms, performing manipulations (e.g. systematic lesions, systematic changes in neuronal morphology or in synaptic transmission) and making systematic measurements (e.g. measurements of cell activity and synaptic dynamics), impossible in animals or in human subjects. These methods should make it possible to obtain new insights into the neuronal circuitry responsible for the model's capabilities, confirming or disconfirming theoretical hypotheses, and guiding the development of technologies inspired by these insights.
6. Where appropriate, the trained brain model would be exported to digital neuromorphic devices allowing physical robots to perform the experimental task in real time, in a physical environment. Such physical robots would provide a starting point for future applications (see below).

Pilot studies and open calls should encourage experimental investigations of a broad range of perceptual, cognitive and motor capabilities, beginning with capabilities that are relatively simple and gradually moving towards more advanced functionality. Candidate capabilities could include basic visual, auditory and somatosensory processing including multisensory perception; object recognition (recognition of faces, body parts, houses, words etc.); action recognition; novelty detection (e.g. auditory novelty detection through mismatch negativity); motivation, emotion and reward; premotor transformations, motor planning and execution of motor behaviour; representations of the spatial environment and navigation; decision-making and error correction; information maintenance and memory encoding; working memory, time-dependent stabilization of cortical representations; and language production and processing.

Accelerated Neuroscience

- Tools for massive data management
- Internet accessible collaborative tools
- Brain atlases and encyclopedia
- Data intensive computing tools
- Data and knowledge predictors
- *In silico* systems for experiments
- Closed-loop technology
- Theory for bridging scales
- Multi-level view of brain function

Using HBP capabilities to understand, diagnose and treat neurological and psychiatric disease

Objective

The HBP seeks to accelerate research into the causes, diagnosis and treatment of neurological and psychiatric disease (see Figure 29). As a first step, the HBP should use the *Medical Informatics Platform* and the data it generates to identify biological signatures for specific disease processes, at different levels of biological organisation. This work would lead towards a new nosological classification based on predisposing factors and biological dysfunctions rather than symptoms and syndromes. We propose that pilot projects should test this strategy for autism, depression and Alzheimer's disease. Open calls for proposals would encourage outside researchers to extend this work and to investigate other diseases.

The second goal should be to use biological signatures of disease as a source of insights into disease processes and to use modelling and simulation as tools to investigate hypotheses of disease causation.

The third goal should be to use disease models to identify potential drug targets and other possible treatment strategies and to simulate their desirable and potentially adverse effects.

The fourth goal should be to develop strategies for personalised medicine, allowing the development of treatment strategies adapted to the specific condition of individual patients.

State of the art

Presently there are very few neurological and psychiatric diseases whose causes are fully understood even when their patho-anatomy and patho-physiology are largely known. For example, in Parkinson's disease we still do not understand the steps that lead from degeneration of less than a million specific nigro-striatal cells to the first clinical symptoms (tremor, akinesia), which only appear when 60% of these cells have already been lost [103]. In a small proportion of cases, the damage is due to exogenous poisons [124]. In many cases, however, the triggering factor(s) is unknown. This situation is complicated by the fact that other relatively common brain diseases have similar Parkinsonian manifestations. It is not known why such symptoms are so common.

Information from Genome-Wide Association Studies (GWAS) has made it increasingly clear that many diseases with different biological causes (e.g., the spino-cerebellar ataxias and multiple associated mutations) present with similar symptoms and *vice versa* (e.g., Huntington's disease presenting with emotional disorders, cognitive deficits or movement disorder). These relationships make it difficult to identify specific drug targets, and to create homogeneous trial cohorts. These are some of the reasons why many pharmaceutical companies have withdrawn from brain research.

Problems with current systems of disease classification and scientific advances – particularly in genetics – are slowly leading researchers to shift their attention from syndromic to biologically-grounded classifications of disease. Until recently, for instance, the dementias were still diagnosed in

terms of dementing syndromes, which often failed to match final *post mortem* analyses. Today, by contrast, clinicians are beginning to interpret neurodegenerative disorders, including the dementias, as diseases of protein misfolding [125]. The *Medical Informatics Platform* would place Europe in a position in which it could pioneer this new biological approach to nosology.

Another area of research, of great relevance to the HBP, is simulation-based pharmacology. Current applications of simulation in drug design focus on the dynamics of molecular interactions between drugs and their targets. To date however, there has been little or no work simulating the complex cascade of events that determines desirable or adverse effects at higher levels of biological organisation. The lack of

effective methods to predict these effects may be one reason for the high rate of failure of CNS drugs in clinical trials. Recent pharmacogenetic studies of anticonvulsants (patient responsiveness to positive drug effects and predisposition to adverse effects) support this hypothesis [126].

Methodology

Categorising human brain diseases. A first important goal for the HBP should be to identify specific biological signatures that characterise disease processes. The discovery of such signatures would result in a new nosology, based on objective and reproducible biological and clinical data such as brain scans of various types, electrophysiology, electroencephalography, genotyping, metabolic, biochemical and

Use Case 2: Developing new drugs for brain disorders

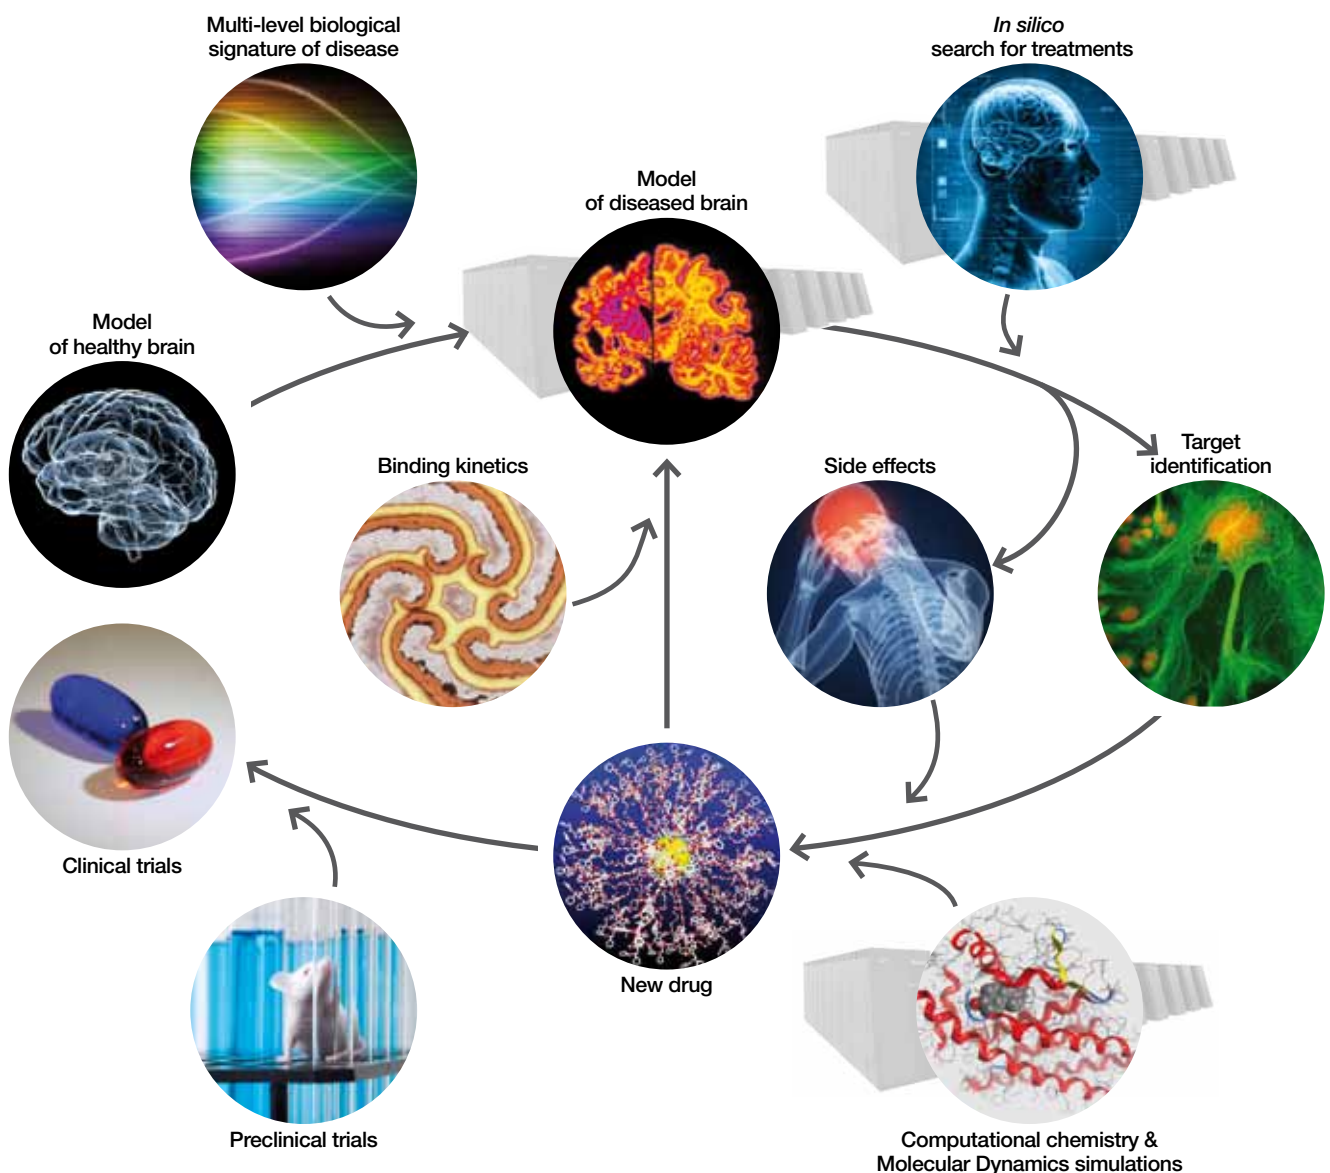


Figure 29: Use of the HBP platforms to accelerate drug development. Optimizing drug discovery and clinical trials

haematological profiles and validated clinical instruments providing quantitative measurements of emotion and behaviour. Initial work by the HBP should focus on the biologically grounded categorisation of autism, depression and Alzheimer's. However, a large part of the overall budget should be reserved for open calls, encouraging research by scientists from outside the Consortium. The calls should encourage systematic study of the full range of neurological and psychiatric disease, making no distinction between disorders of perception, cognition, action, mood, emotion and behaviour.

Simulate hypotheses of disease causation. The discovery of biological signatures for a disease would suggest hypotheses of disease causation. The *Brain Simulation Platform* should therefore allow researchers to model alterations in brain physiology and structure they believe to be implicated in different diseases and to simulate the complex non-linear interactions leading to changes in cognition and behaviour. The realisation that the brain is not susceptible to linear analysis has come slowly.

Again, we are at a tipping point – the *Brain Simulation Platform* would make it possible to simulate the effects of brain lesions on the overall functioning of brain systems, including the short-term, adaptive plasticity effects that normally palliate lesions. Simulation would also facilitate the testing of causative hypotheses for diseases for which there are no available animal models, and for disorders where such models are inadequate, for example, when disorders are associated with defects in higher cognitive function. Simulation would also teach researchers to distinguish between causative and secondary alterations associated with disease processes. The success of this kind of research will be judged by its contribution to understanding the role of different levels of biological organisation in brain disease, and to identifying new targets for treatment.

Simulation-based testing of drugs and other treatments for brain disease. An important goal for the HBP should be to provide tools that allow researchers to simulate the effects of treatments (drugs, non-pharmacological treatments) at different levels of biological organisation. The rules that govern brain organisation, structure and function at multiple spatial and temporal scales, once identified and built into unified

models of the brain, would become targets for (mathematical) modification. The effects of such point disturbances would be identifiable elsewhere and at different levels of brain expression. The identification of the sites of disorganisation, dysfunction or anatomical change coming from the definition of disease signatures would provide the synergies needed for a new method of drug discovery.

Services for personalised medicine

The discovery of reliable biological signatures for psychiatric and neurological disorders would represent a major step towards personalised medicine in which treatments are tailored to the conditions of individual patients. The HBP should collaborate actively with hospitals and industry to develop projects that implement and validate such techniques.

Using HBP capabilities to develop future computing technologies

Objective

One of the main goals of the HBP should be to apply improved understanding of the brain to the development of novel computing technology. The project should use its ICT platforms to promote research projects aimed at developing novel software, hardware and robotic systems, inspired by knowledge of the brain and at exploring their possible applications (see Figure 30). Such technologies would have the potential to overcome critical limitations of current ICT, including limits on programmability, power consumption and reliability. The end result would be novel applications with a potentially revolutionary impact on manufacturing, services, health care, the home, and other sectors of the economy. As in other areas of HBP applications work, the majority of this research should be entrusted to researchers who were not involved in the original HBP project, selected via open calls for proposals.

State of the art

Although the numbers of components per chip and performance per unit of investment continue to grow exponentially, other measures of computing performance such as power consumption per chip or clock speeds on digital processors have already reached saturation. On measures such as the number of components per unit area, current technology is rapidly approaching fundamental limits imposed by the atomic structure of matter. Already, today's deep-submicron technologies suffer from extreme requirements in lithography or production technology, making investment in chip foundries a multi-billion dollar endeavour. These trends go hand in hand with ever increasing software complexity. In particular, modern computer processors run and communicate continuously, consuming large quantities of power. The computer industry is already exploiting parallelism and redundancy in many-core processors and many-processor computing systems interconnected by high bandwidth and low-latency interconnection fabrics. Brain-inspired technologies can take this parallelism to the extreme, opening the road to low-power, highly reliable systems with brain-like intelligence [2].

Accelerated Medicine

- Tools for massive data management
- Informatics tools for federated data mining
- Informatics tools to derive biological signatures of disease
- Diagnostic tools
- Multi-level models of brain disease
- Simulations for hypotheses of disease causation
- *In silico* testing of treatments
- *In silico* drug development

Use Case 3: Developing neuromorphic controllers for car engines

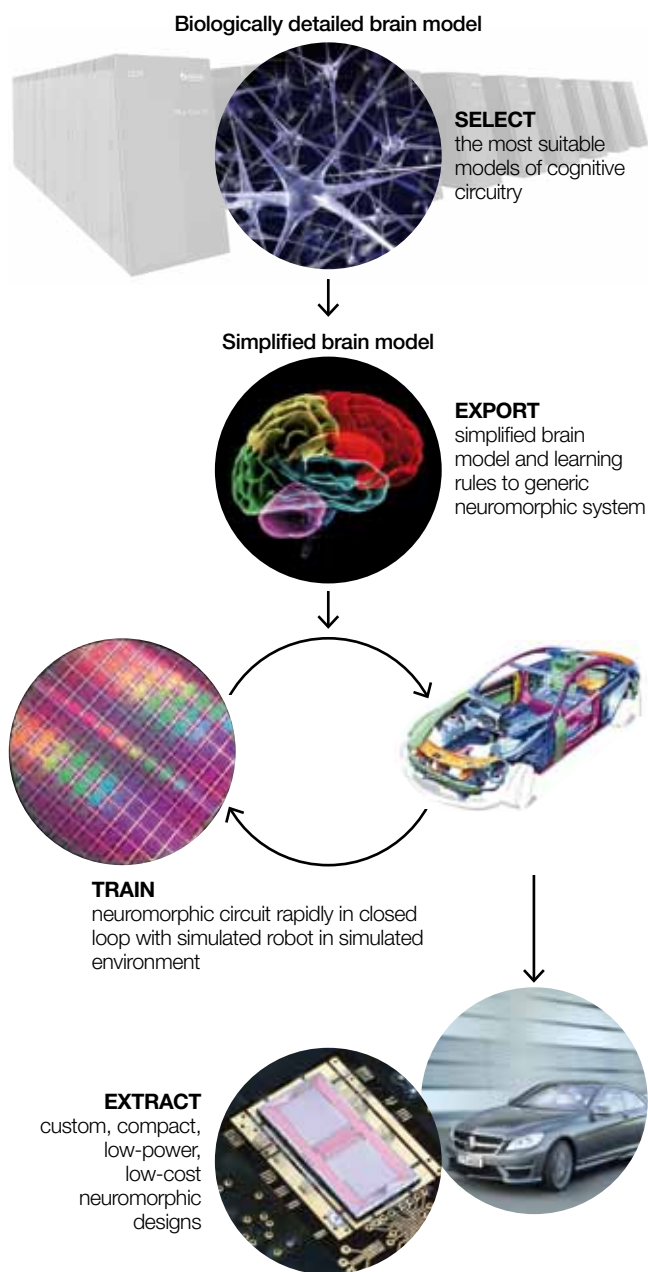


Figure 30: Use of the HBP platforms to accelerate the development of future computing technologies. Four steps: brain simulation; simplification; rapid prototyping; low-cost, low-power chips

Methodology

The HBP should collaborate with industry partners and researchers from outside the HBP Consortium to demonstrate the potential of the project’s ICT platforms for the development of novel ICT systems and applications, inspired by the architecture of the brain. The HBP Consortium should consider proposals in any area of systems or applications development. The main areas of interest should include the following.

Accelerated Future Computing

- Interactive supercomputing
- Massive data mangement
- Neuromorphic enhanced
 - Computing
 - Mobile
 - Entertainment
 - Industrial applications
 - Vehicles
 - Datamining
 - Home devices
 - Robots

High performance computing: neuromorphic cores for conventional high performance computers, brain-inspired communications protocols; brain-inspired strategies for information storage and retrieval; massively parallel very low-power computing cores.

Software: applications incorporating advanced capabilities for pattern recognition, feature recognition, motor control, decision-making etc. (e.g. applications in industrial control, image processing, language processing etc.).

Neuromorphic computing systems and devices: neuromorphic controllers for manufacturing, household appliances, vehicles, image and video processing, mobile telecommunications etc.; neuromorphic processors for use in high performance computing, neuromorphic processors for use in commodity computers and mobile devices.

Robotics: specialised neurorobotic systems for applications in manufacturing, services, health-care, ambient assisted living, the home and entertainment.

The HBP Society and Ethics Programme

Objectives

As just described, brain simulation and the technologies to which it could give rise have numerous social, ethical and philosophical implications. Stakeholders thus have an interest in recognising concerns early and in addressing them in an open and transparent manner. In particular, early engagement can provide scientists with opportunities to gauge public reaction to their work, and to hone their research objectives and processes in the light of these reactions. We therefore recommend that the HBP launch a major Society and Ethics Programme. The goal of the programme should be to explore the project’s social, ethical and philosophical implications, promoting engagement with decision-makers and the general public, raising social and ethical awareness among project participants, and ensuring that the project is governed in a way that ensures full compliance with relevant legal and ethical norms. The programme should draw on the methods developed during empirical investigations of emerging technologies in genomics, neuroscience, synthetic biology, nanotechnology and information and communication technologies [127] as well as on the biomedical tradition of engaging with ethical issues through the application of

formal principles [128] – now usually implemented through ethical review processes.

State of the art

Social consequences of the HBP

HBP research entails high expectations of social and economic benefits. However, the impact of basic research results on society often depends not so much on the research itself as on developments in apparently unconnected areas of science and technology or on social, political and legal factors external to science [129-131].

Current approaches to forecasting development pathways use one of two strategies. The first studies the views, attitudes and strategies of key stakeholders with methods from the empirical social sciences such as laboratory ethnographies [132, 133]; the second, which has reached its highest stage of development in the UK (www.bis.gov.uk/foresight), uses systematic foresight techniques such as modelling, horizon scanning and scenario planning. Crucially, this kind of study always includes an assessment of key ethical concerns such as privacy, autonomy, transparency, the appropriate balance of risks and benefits, responsibility and accountability, equity and justice [134].

Conceptual and philosophical issues

Since the 1960s, scientific and technical advances [135] have made it ever easier to anatomise the brain at the molecular, cellular and circuit levels, encouraging claims that neuroscience is close to identifying the physical basis of mind. Such claims have major implications not only for medicine but also for policies and practices dealing with normal and abnormal human conduct, and for conceptions of personhood. The significance and consequences of these developments are strongly debated, with some authors arguing that we now know enough to understand the neural bases of human selfhood and higher mental functions [136, 137], while for others, the *neuroreductionist* model attributes capacities to brains that can only properly be attributed to persons [138, 139]. Some have suggested that progress in neuroscience will lead to radical improvements in our ability to treat psychiatric disease [140, 141]; others are more doubtful [142, 143]. While functional imaging has been crucial in the development of new conceptualisations of human mental states, many leading researchers are highly critical [144].

Meanwhile, studies of the neural basis of higher brain functions have fed scientific and semi-popular debates about ideas of personhood [145-147] and free will [148-150] while studies combining psychophysics and brain imaging (e.g., [151]) have encouraged philosophers to readdress the eternal mystery of conscious awareness. The emerging discipline of *neuroethics*, a home for some of these discussions, has produced an extensive literature both on general conceptual issues [152-155], and on specific questions such as the functional neuroimaging of individuals belonging to different ethnic and age groups [156, 157], cognitive enhancement and memory distortion [158-161], neuroscience and law [162-164] and potential military applications of neuroscience [165]. The capabilities developed by the HBP would provide new material for these debates.

The public, dialogue and engagement

Attempts to achieve public dialogue and engagement during the development of new technologies [166, 167] have used a range of methods and approaches [168] including consensus conferences, citizen juries, stakeholder workshops, deliberative polling, focus groups and various forms of public dialogue. In the UK, for example, the dialogue is organised through the nationally funded ‘ScienceWise’ initiative (see <http://www.sciencewise-erc.org.uk/>). Other notable work in this area has been undertaken by the Rathenau Institute in the Netherlands (<http://www.rathenau.nl/en.html>). The motivations for such exercises [130, 169, 170] are sometimes normative – it is easy to argue that citizens affected by research have a right to participate in crucial decision-making – sometimes instrumental. Many authors have argued, for instance, that dialogue can reduce conflict, help to build trust and smooth the introduction of innovative technology. The strongest conclusion from these debates is that not even the best prepared exercises can comprehensively represent the positions of all parts of society or resolve the issue of which groups or opinions should be given most weight in a decision. It is important, therefore, that such exercises should respect scientists’ legitimate desire to inform the public about their research, while avoiding self-conscious attempts to steer public opinion in a particular direction. Experience from other areas of emerging technology research shows that this requires a sensitive approach [171]. Public engagement exercises are successful only if participants are convinced that they can genuinely influence the course of events [172].

Researcher awareness

Ethical issues cannot be reduced to simple algorithms or prescriptions: moral statements and positions always require higher-level ethical reflection and justification. From an ethical point of view, this reflection should come, not just from external “ethical experts”, but also from researchers and their leaders. This kind of general reflexivity is currently not the norm and is likely to meet resistance. Studies suggest that the best way of achieving it is to embed measures to raise researcher awareness in governance structures [173], a technique already applied in other areas of cutting-edge technical research, notably nanotechnology (www.nanocode.eu) [174] and synthetic biology.

Governance and regulation

Today’s science regulatory environment is a result of research that provoked a vigorous social and governmental response [175]. One example is animal research, in which the response took the form of The Council of Europe’s Convention for the Protection of Vertebrate Animals used for Experimental and other Scientific Purposes (ETS 123) (1985), and the EU Directive for the Protection of Vertebrate Animals used for Experimental and other Scientific Purposes [176] – documents that have set European standards for the use of mice and other vertebrates in the laboratory. Another more recent example is the case of synthetic biology. In this case, the reaction came only after a private institution had created the first self-replicating bacterial cell from a com-

pletely synthetic genome [127]. Equally compelling cases can be gleaned from biomedicine, genetics, information and computer technology, bioengineering, neurorobotics, and nanotechnology [171].

Modern governance of innovation in biotechnology involves a variety of actors, including research organisations, national and supranational regulators, governmental or quasi-governmental organisations, professional bodies, publishers of science journals, and representatives of the mass media and public opinion. As Gottweis [177] noted for the case of transnational research on embryonic stem cells, decision-making takes place "... at the fuzzy intersection between science, society, and politics". This is complicated, in the case of international projects, by the need to take account of different national jurisdictions.

Methodology

The Human Genome Project's Ethical, Legal and Social Issues (ELSI) programme [178] – which absorbed 3-5% of the project's total budget – demonstrated that open public discussion is an effective strategy for handling potentially controversial issues raised by scientific research. We recommend that the HBP should learn the lessons of this programme, setting up its own Society and Ethics Programme, running for the whole duration of the project. The programme would bring together scholars in the brain sciences, social sciences, and the humanities to study and discuss relevant issues, using all available channels to encourage open, well-informed public debate.

As a contribution to these goals, the HBP should organise a detailed programme of foresight exercises and of academic research into the project's social, economic, legal, and philosophical impacts and their ethical implications. This programme should be accompanied by specific measures – a European Citizen's Deliberation, citizen juries, consensus conferences, web-based dialogue tools, education programmes – that previous projects have shown to be effective. There should also be a parallel programme to raise awareness of social and ethical issues *within* the HBP Consortium. Finally the HBP should design a detailed system of ethical governance to ensure that research carried out within the project meets the highest possible ethical standards and that it complies with relevant law and regulations. The governance system should include an *Ethical, Legal and Social Aspects Committee* that oversees the overall activities of the project and a *Research Ethics Committee* that collects and reviews HBP research ethics applications prior to submission to external Independent Review Boards.

Coordination of resources and research communities

One of the Human Brain Project's most important goals should be to catalyse integration between different areas of neuroscience, medicine and ICT. The HBP-PS started the process by bringing together more than 300 researchers in a broad range of disciplines, from cellular and molecular level neuro-

science to neuromorphic computing and ethics. The project led to highly productive exchanges of ideas between scientists who would not meet in the normal course of academic affairs, for example between electrophysiologists and cognitive neuroscientists, cognitive neuroscientists and roboticists, roboticists and specialists in neuromorphic computing.

The Human Brain Project and its ICT platforms would place this kind of collaboration on new foundations (see Table 1).

1. The HBP would be a mission-oriented project with a small number of well-defined goals. The joint effort to achieve these goals would create large, new incentives for interdisciplinary collaboration.
2. The HBP ICT Platforms would provide scientists with access to advanced technical capabilities (brain atlases, brain simulation, high performance computing, medical informatics, neuromorphic computing, neurorobotics) whose complexity and cost preclude their use by all but the best-funded research groups. Access would be open, not just to scientists within the HBP Consortium but to the entire international scientific community. The availability of these resources has the potential to revolutionise current research practices, which are currently constrained by the resources and know-how available to individual groups.
3. The tools and data made available by the *Neuroinformatics*, *Medical Informatics* and *Brain Simulation Platforms* (brain atlases, federated clinical data, unifying models) would encourage data sharing among groups studying different aspects of the brain and different levels of brain organisation. This represents a potentially revolutionary change in current research practices. Unifying models and simulations of the brain, would make it easier for groups in different areas of neuroscience to interpret their results in the context of data from other groups. This again represents a major change.
4. Developing applications in neuroscience, medicine and future computing technology would provide strong incentives to collaboration between scientists from different disciplines. To cite just one example, the design and implementation of closed-loop experiments in cognitive neuroscience would involve collaboration between experimental and theoretical neuroscientists, as well as experts in brain simulation, high performance computing, neuromorphic computing and neurorobotics.
5. Calls for proposals, research grants and studentships funded by the HBP (see p. 70) would encourage the use of the project's capabilities by scientists from outside the initial HBP Consortium, especially young investigators.
6. The very existence of the Human Brain Project would provide an organisational framework for frequent meetings and exchanges among scientists from different research communities. Many kinds of collaboration foreseen by the project (e.g. between molecular and cognitive neuroscience, between high performance computing and neuroscience, between neuroscience and robotics, between robotics and neuromorphic computing) are currently very rare or completely absent.

Activity	Uses results from	Observations
Revealing integrative principles of cognition	Multi-level structure of the mouse brain Multi-level structure of the human brain Brain function and cognitive architectures Theory All ICT platforms	Experimental design would be guided by data and protocols from cognitive neuroscience. Data for the project would come from all areas of neuroscience present in the project, and would include clinical data generated by the <i>Medical Informatics Platform</i> . Preliminary processing of the data and predictive modelling would use the capabilities of the <i>Neuroinformatics Platform</i> . Experiments would use brain models developed with the <i>Brain Simulation Platform</i> . The <i>Neurorobotics Platform</i> would make it possible to design and execute closed-loop experiments. The <i>High Performance Computing Platform</i> would provide the computing power to run experiments using detailed brain models. The <i>Neuromorphic Computing Platform</i> would provide hardware support for simpler models running in real time or many times faster (for experiments involving learning over long periods of simulated time).
Understanding, diagnosing and treating disease	Multi-level structure of the mouse brain Multi-level structure of the human brain Brain function and cognitive architectures Theory All ICT platforms	The <i>Medical Informatics Platform</i> would provide clinical data and the software to analyse the data. This work would make it possible to identify biological signatures of disease and to formulate hypotheses of disease causation. Such hypotheses would be informed by basic neuroscience data generated elsewhere in the project. The other platforms would make it possible to test them in <i>in silico</i> experiments, identifying potential targets for treatment, and screening candidate drugs.
Developing future computing technologies	Theory, <i>Brain Simulation Platform</i> , <i>High Performance Computing Platform</i> , <i>Neuromorphic Computing Platform</i> , <i>Neurorobotics Platform</i>	HBP work in theoretical neuroscience would make it possible to simplify the detailed brain models produced by the <i>Brain Simulation Platform</i> . The <i>Neuromorphic Computing Platform</i> would allow researchers to explore the potential of systems and devices based on such circuits. The <i>Neurorobotics Platform</i> would provide the tools to explore applications in robotics. External research groups would be encouraged to explore applications in high performance computing and for other kinds of generic software.

Table 1: Use of the ICT platforms

4 Implementation

The challenge

The HBP would be a large, interdisciplinary ten-year project with an overall budget of more than Euro 1 billion and a large number of partners in many different countries. This makes it essential to define management and governance models that allow effective coordination between the Consortium, the European Commission and national funding agencies and within the Consortium itself. The models chosen should meet three key requirements.

- They should support the strong leadership and tight project management necessary for the HBP to achieve its scientific and technical goals.
- They should guarantee that the ICT platforms become a genuine resource for the scientific community – ensuring technical quality and performance as well as 24/7 availability.
- They should guarantee the openness and flexibility of the project, ensuring that it draws the maximum possible benefit from new ideas and new technologies as they emerge, and encouraging the broadest possible participation by researchers with different theoretical and experimental approaches.

Research organisation

Different classes of research

Core research activities (data generation, building and operating the platforms, pilot projects demonstrating the value of the platforms) should be based on a detailed work plan, which provides precise definitions of the work to be delivered, milestones, costs, and the partners responsible. In cases where the necessary competences are not present in the HBP Consortium the work should be assigned by competitive call (see below). HBP work plans should be revised on a regular basis. We propose revisions at the end of the ramp-up phase, in Year five and in Year seven.

Research using the ICT platforms should take the form of research projects, selected through a competitive process open to the entire scientific community. The process should be modelled on the current Marie-Curie, ERC and FET

*The Human Brain Project
will require an international effort
and an unprecedented scale
of multidisciplinary collaboration*

Programmes. Priorities should be defined in work programmes agreed between the HBP Consortium, the European Commission and national funding agencies. As in the FET Programme, the work programmes would define funding, themes and a schedule for calls for proposals (on average one per year). Researchers responding to the calls would formulate research proposals, which would be evaluated by experts external to the HBP Consortium. Selected projects would draw on results from many different areas of HBP research as shown in the table overleaf.

Education

The HBP should support its research activities with large-scale programmes of multidisciplinary education for scientists inside and outside the HBP Consortium. These should include an HBP Graduate School and programmes providing training in the use of the HBP platforms. The programme would be coordinated by central management.

Divisions and work packages

To facilitate the coordination of HBP research, HBP researchers should be organised into divisions, each dedicated to a specific discipline, as shown in Figure 31. Each division would contribute to multiple scientific activities. For example, the High Performance Computing Division would build and operate the *High Performance Computing Platform*, but would also contribute to Neuroinformatics (cloud computing), Brain Simulation (computational support), Medical Informatics (cloud computing), Neuromorphic Computing (design of neuromorphic devices and systems) as well as to applications development.

Each of the planned activities would be broken down into work packages, comparable in size to an Integrated Project in the Framework 7 Programme.

Organisation of HBP Scientific & Technological Work		Divisions								
		Molecular and Cellular Neuroscience	Cognitive Neuroscience	Theoretical Neuroscience	Neuroinformatics	Brain Simulation	High Performance Computing	Medical Informatics	Neuromorphic Computing	Neurorobotics
Scientific and technological research	Mouse Data	Light	Light	Light						
	Human Data	Light	Light	Light						
	Cognition Data	Light	Light	Light						
	Theory	Light	Light	Dark	Light	Light	Light	Light	Light	Light
	NIP			Light	Dark	Light	Light	Light	Light	Light
	BSP			Light	Light	Dark	Light	Light	Light	Light
	HPCP			Light	Light	Light	Dark	Light	Light	Light
	MIP			Light	Light	Light	Light	Dark	Light	Light
	NMCP			Light	Light	Light	Light	Light	Dark	Light
	NRP			Light	Light	Light	Light	Light	Light	Dark
	Applications for neuroscience	Light	Light	Light	Light	Light	Light	Light	Light	Dark
	Applications for medicine	Light	Light	Light	Light	Light	Light	Light	Light	Dark
	Applications for future computing	Light	Light	Light	Light	Light	Light	Light	Light	Dark

Figure 31: Organisation of work in the HBP. The nine divisions contribute to data generation, platform building and platform use. Theory is a cross-cutting activity, present in all HBP activities. Dark shaded areas represent responsibilities for work packages; light shaded areas show contributions to the work programme

Governance

The initial funding for the HBP would be provided under the Framework 7 Programme, and would be based on Framework 7 administrative and financial rules. In the first phase of the project, the governance of the project would follow the conventional model used for FP7 Cooperative Projects. In this model, the execution of the project would be the responsibility of a consortium of partners (the *HBP Consortium*), whose composition would evolve as new groups join the project and old ones leave. Legal arrangements among the partners (governance bodies and procedures, arrangements for the admission of new partners, management of IP etc.) would be specified in a *Consortium Agreement*. The Grant Agreement with the EU and the terms of the *Consortium Agreement* would be designed to facilitate such changes with a minimum of administrative and legal overhead. By the terms of the *Consortium Agree-*

ment, overall coordination would be entrusted to a *Coordinating Partner*, which would provide a unique point of contact for the European Commission, coordinating the activities of the other partners, each of which would have an independent contractual relationship with the EU.

Governing bodies

The supreme governing body for the project would be an *HBP Assembly*, which would include representatives of all partners. Each representative would have a number of votes proportional to the partner’s share in the overall HBP budget. The *HBP Assembly* would normally meet once a year.

Strategic management decisions would be the responsibility of a ten to twelve member *HBP Governing Board*, elected by the *Assembly*. The *Governing Board*, which would include a minimum of three non-executive directors from outside the Consortium, would have a five-year mandate, which the *Assembly* could revoke at any time. As far as

possible, decisions would be taken by consensus. Each member of the *Governing Board* would have one vote. Ordinary decision-making would be by simple majority vote. Some special kinds of decision, defined in the *Consortium Agreement*, would require a larger majority. The *Governing Board* would meet at least every three months.

The *Governing Board* would be assisted by an independent *Science Advisory Board*, which would monitor the project and provide scientific advice. The *Science Advisory Board* would not have decision-making powers.

Further assistance would be provided by the *Presidents' Council*, made up of the presidents of the universities and

The Human Brain Project will require a professional management and administration team

research institutions participating in the project. The *Presidents' Council* would help to ensure effective alignment between HBP and partner strategies.

Governance and management organigramme

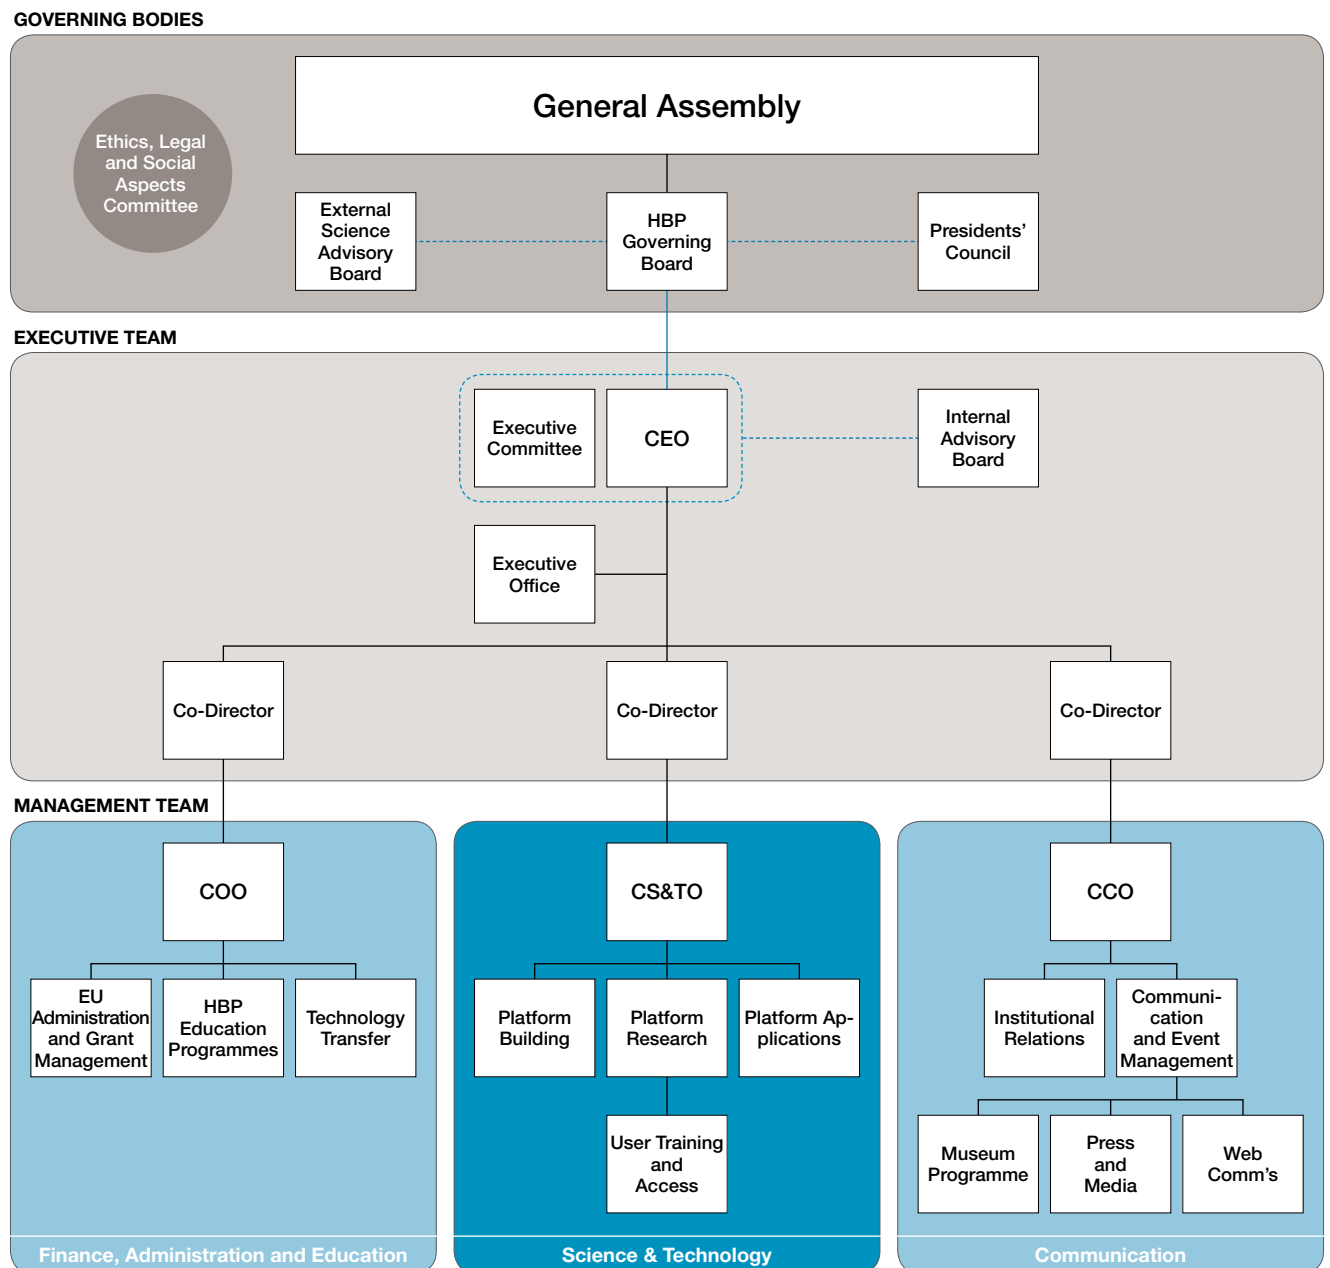


Figure 32: Governance and management: The HBP Board supervises the Executive Team, and is advised by the Science Advisory Board and the Presidents' Council. The Executive Team is responsible for the timely execution and development of the project work plan. The Management Team is responsible for the day-to-day operation of the project

Ethics

Strategic oversight over ethical, legal and social issues would be provided by an independent *Ethical, Legal and Social Aspects Committee*, made up of scientists, legal experts and lay representatives from outside the project. The composition of the committee would be defined to ensure adequate representation for different viewpoints and opinions and different approaches to ethical and social issues. The committee would provide advice to the *HBP Assembly* and the *HBP Governing Board* on its own initiative and in response to specific requests. A separate *Research Ethics Committee* would be responsible for day-to-day governance of ethical issues, including the collection and review of HBP ethics applications prior to their submission to external Independent Review Boards.

Executive Team

The *Board* would elect a *Chief Executive Officer (CEO)* and an *Executive Committee* comprising three *co-directors* and the *CEO*. The *Executive Committee* would be responsible for executing the decisions of the *Board* and managing the project. The *CEO* and the *Executive Committee* would both have a five-year mandate, which the *Board* could revoke at any time.

The *CEO* and the *Executive Committee* would be supported by an *Internal Advisory Board*, with responsibility for mediating conflicts between the partners, and an *Executive Office*, providing administrative support.

The *CEO* and the *Executive Committee* would be assisted by a permanent *Operations Team*, as described below.

Possible modifications of the governance model for Horizon 2020

Previous experience in the Framework 7 Programme suggests that the standard governance model for EU Cooperative Projects may not meet all the requirements of a very large, interdisciplinary project such as the HBP. The main weakness of this model is that it requires a contractual relationship between each individual partner and the EU. It would thus be desirable to create a new model that simplifies these relationships. Mechanisms would depend on the administrative and financial rules for Horizon 2020, which are not yet available. As input to the decision-making process, we suggest the following.

1. Once the Horizon 2020 Programme is in place, the HBP would transition to a new model of governance in which the *Coordinating Partner* would be an independent legal entity (the *HBP International Association*), jointly owned and governed by the other partners in the project (see below).
2. The universities, research institutions and other member organisations participating in the original *HBP Consortium* would organise themselves into national groupings (*National Associations*), which become beneficiaries of the EU Grant Agreement. The *National Associations* would be empowered to apply for and coordinate funding

Governance model for Horizon 2020

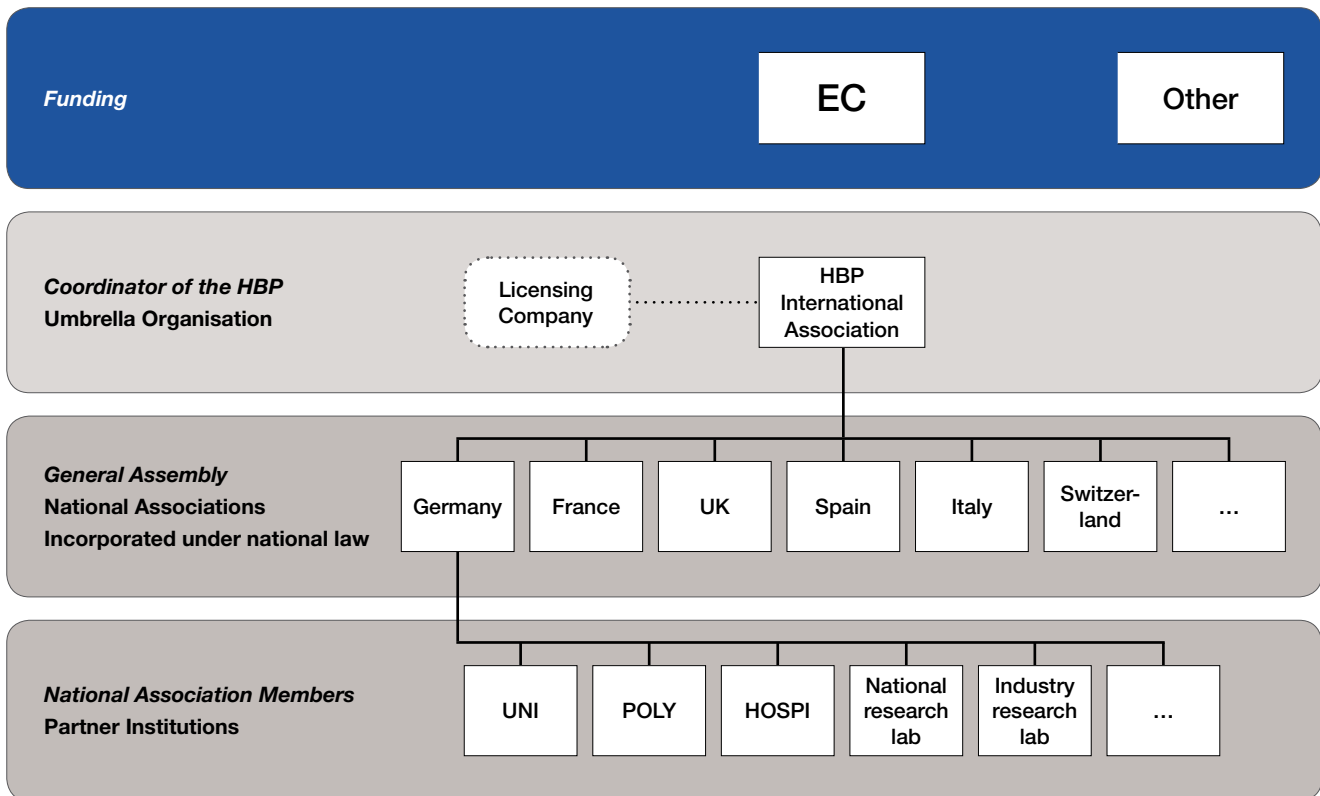


Figure 33: Proposed governance model for the HBP as a FET Flagship in the Horizon 2020 Programme

from national funding agencies, which they would distribute to universities and research institutions within their own countries. Each of the institutions participating in HBP research work would be a member of the appropriate *National Association* and would have equal voting rights within the *Association*. When new institutions joined the project, they would join the relevant *National Association*.

3. All other aspects of governance would remain the same as in the ramp-up phase.

Figure 33 illustrates this scheme. The actual governance structure for the operational phase would depend on Commission rules and policies yet to be defined. The HBP Consortium would define its final proposals once these rules and policies are known.

Management

Operations Team

The HBP *Executive Team* should be assisted by an *Operations Team* made up of three departments, each headed by a professional manager with a permanent staff, as shown in Figure 32.

Finance, administration and education

This department should be responsible for all issues related to HBP financial and administrative affairs and to the project's education programme. These would include coordination with the EU administrative, legal and financial departments, and coordination with financial and administrative officers in the partner organisations, ensuring compliance with EU and project reporting requirements, grants management (calls for proposals, contract negotiation, progress monitoring, review management, payments), and monitoring and control of financial resources. The department would also be responsible for managing the project's education programme, including long-term programmes of education, training events, training in the use of the HBP platforms and distance learning services. A special office would be responsible for technology transfer (relations with industry, IP management, organisation and support for spin-off companies).

Science and technology

This department would be responsible for coordinating all science and technology work within the project (data, theory, platforms and applications), ensuring that work is delivered on time and on budget and that it is properly reviewed. It is this department that would coordinate the building and operation of the ICT platforms, ensuring that they provide a high quality of service on a continuous basis. The department would work with teams in the partner organisations to design and enforce common software engineering methodologies, common development platforms and programming standards, common verification and validation procedures, common standards for documentation, training and techni-

cal support, common human computer interfaces, common interfaces between platforms, common procedures for data security, common service level agreements and common procedures to monitor their application.

Communications

The communications department should coordinate all central management and partner activities related to communications inside and outside the Consortium. These include communications with external institutions (international organisations, EU and Member State institutions, research projects in related areas, NGOs etc.), communications with the media and online communications, and event management (public events, project meetings, shared videoconferencing facilities). The department's responsibilities would include the organisation of a long-term programme of collaboration with European science museums.

Costs and revenues

Overview

Our estimated total cost of the HBP is **Eur 1,190 million**, spread over a period of ten years (see Table 2). This comprises a two and a half year ramp up phase (**Eur 80 million**), a four and a half year operational phase (**Eur 673 million**), and a three-year final phase in which the project moves towards financial self-sustainability (**Eur 437 million**). All costs are estimated at fixed 2012 prices and exchange rates.

Of total estimated HBP costs, **Eur 316 million** (27%) would be for data generation, **Eur 456 million** (40%) for developing and operating the ICT platforms, **Eur 221 million** (19%) for applications, **Eur 96 million** (8%) for theoretical neuroscience and the HBP Society and Ethics Programme and **Eur 71 million** (6%) for management (see Figure 34). An additional **Eur 30 million** would go to projects funded by a future ERANET+ programme (see below).

Eur 934 million (78.5%) would go to research activities carried out by the initial HBP Consortium (see Figure 35). **Eur 256 million** (21.5%) would be dedicated to open calls. As shown in Figure 37, the relative amount of funding dedicated to open calls would increase steadily over the duration of the project. By the end of the project this funding would account for 40% of the budget. The rest of the budget would be dedicated to updating and maintaining the platforms.

The HBP's main sources of revenue should consist of:

- Research financing from the European FP7 and Horizon 2020 Programmes
- Co-funding by Consortium partners
- Funding under a future ERANET or ERANET+ programme

According to our estimates, 54% of the total funding (**Eur 643 million**) would be provided by the European Commission. The remaining 46% (**Eur 547 million**) would come from the Consortium partners and the ERANET or ERANET plus programme. Once the platforms were

Costs by category of activity (core project)




Figure 34: Distribution of costs by major areas of project activity

Costs by division (core project)




Figure 35: Distribution of costs by division

operational and had demonstrated their value, we envisage that public and industry sources would provide additional funding to cover the cost of research projects, going beyond the original goals of the project.

In the ramp-up phase (under FP7), 73% of total funding (Eur 59 million) would come from the European Commission and the remaining 27% (Eur 22 million) from the partners and the ERANET or ERANET+ . In the operational and sustainability phases (under Horizon 2020), the proportion of funding coming from the Commission would fall and the proportion of funding from the partners would rise (Commission Eur 576 million – 53%; partners Eur 510 million – 47%) (see Figure 36).

Costs	€ millions	%
Core Project direct costs	725	60.9
of which		
- Personnel	555	46.6
- Equipment	71	6.0
- Consumables	99	8.3
Overheads (indirects costs = 60% of direct costs)	435	36.6
ERANET+	30	2.5
TOTAL	1,190	100

Table 2: HBP budget by major cost category

Detailed analysis of costs for the core² project

Overview

Table 2 shows the estimated division of costs by major cost category. Table 3 provides a disaggregated view of costs by division. In what follows, we describe these costs in greater detail.

Personnel

Total personnel and personnel-related costs are estimated at **Eur 555 million** (48% of the budget). Table 4 provides a breakdown of the effort covered by these costs, based on an average cost of Eur 75,000 per person-year. The divisions contributing the most effort would be High Performance Computing (1032 person years), Cognitive Neuroscience (996 person years), Neuroinformatics (887 person years) and Brain Simulation (803 person years).

Equipment

Estimated equipment costs for the HBP amount to **Eur 71 million** (6% of the total budget). A large portion of this figure would be accounted for by the project’s contribution to the cost of supercomputing equipment in Jülich and Lugano. It is expected that supercomputing would receive significant additional support from national funding agencies (see below).

Consumables

Estimated consumable costs for the HBP amount to **Eur 99 million** (8% of the budget). More than a third of this expenditure would be for Neuromorphic Computing, which would spend heavily on chip fabrication services.

Indirect costs

Indirect costs are estimated at **Eur 435 million** (37% of the budget). This estimate is based on the assumption that indirect costs would be computed as 60% of direct costs – the *transitional flat rate* used by the majority of universities and research institutions in FP7 projects.

² These costs do not include funding for group proposals under a possible ERANET or ERANET+ programme.

Costs by year

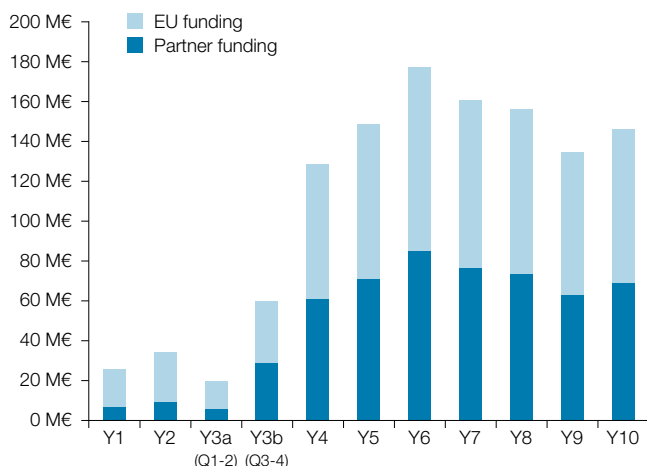


Figure 36: EU and partner contributions to HBP costs over the ten years of the project

Open calls as % of funding




Figure 37: Costs for open calls (support for researchers and projects) rising to 40% of the budget by the end of the project

Division (costs in € millions)	Personnel	%	Equipment	%	Consumables	%	Other & Indirect	%	Total Budget	%
Molecular & Cellular Neuroscience	48.5	9.0%	2.9	4.1%	5.8	5.9%	36.2	8%	93.4	8.1%
Cognitive Neuroscience	74.7	13.9%	4.5	6.3%	9.0	9.1%	55.7	12%	143.9	12.4%
Theoretical Neuroscience	41.9	7.8%	0.4	0.6%	0.2	0.2%	29.2	6%	71.8	6.2%
Medical Informatics	40.3	7.5%	3.6	5.0%	0.0	0.0%	27.8	6%	71.8	6.2%
Neuroinformatics	43.8	8.2%	3.9	5.5%	0.0	0.0%	30.6	7%	78.3	6.7%
Brain Simulation	60.2	11.2%	5.5	7.7%	0.0	0.0%	44.7	10%	110.5	9.5%
High Performance Computing	77.4	14.4%	45.0*	62.9%	6.6	6.6%	80.7	18%	209.7	18.1%
Neuromorphic Computing	25.6	4.8%	0.0	0.0%	70.5	71.5%	60.4	13%	156.5	13.5%
Neurorobotics	66.5	12.4%	5.6	7.8%	6.6	6.6%	50.4	11%	129.0	11.1%
Society & Ethics	14.7	2.7%	0.0	0.0%	0.0	0.0%	9.8	2%	24.5	2.1%
Management	42.4	7.9%	0.0	0.0%	0.0	0.0%	28.1	6%	70.5	6.1%
TOTAL	536.1	100%	71.5	100%	98.7	100%	453.6	100%	1'159.8	100%

Table 3: HBP direct costs by major cost category and by division

* The majority of funding for high performance computing equipment will be provided by national sources

Division	Person-Years
Molecular & Cellular Neuroscience	646
Cognitive Neuroscience	996
Theoretical Neuroscience	559
Medical Informatics	538
Neuroinformatics	584
Brain Simulation	803
High Performance Computing	1'032
Neuromorphic Computing	341
Neurorobotics	887
Society & Ethics	196
Management	566
TOTAL	7'148

Table 4: HBP: personnel effort per division for the full duration of the project (person years)

Sustainability

The HBP should ensure that the project continues to offer its capabilities to scientists and to address key scientific challenges, after the initial 10-year period. This would require new sources of revenue that complement and eventually replace the original funding. Possible sources of funding include:

- Academic and commercial organisations outside the HBP that contribute to the capital and running costs of the facilities in return for guaranteed access³
- Commercial research projects paid for by industry (primarily pharmaceutical companies, computer manufacturers, manufacturers of medical devices)

³ This is similar to the funding model for large telescopes where national funding agencies cover a percentage of the capital and running costs and reserve a percentage of observation for scientists from a given country.

- Licensing of IP: the project would license IP to interested third parties through the licensing company described on page 81
- Direct industry funding of research: the HBP would collaborate with industry to develop new intellectual property. In these cases industry would fund staff, equipment and consumables and contribute to the cost of the platforms.

The value of this funding would depend on the actual development of the project. We expect that in the last years of the project it would enable the HBP to pursue applications work not explicitly included in the original work plan.

Accountability

As a project with a very large budget, and many partners, it is essential that the Human Brain Project be fully accountable to the European Commission, the Member States and above all to the taxpayers who would fund HBP research. At the formal level, the project should take measures to ensure that its annual financial statements and financial statements from partners receiving substantial funding are certified by independent auditors. In addition, the project should also organise internal audits of randomly selected partners and activities, ensuring that all partners are audited at least every 3-4 years.

On a less formal level, the HBP should adopt a policy of maximum openness ensuring that its research facilities are open for visits by representatives of European and national institutions and the media, and ensuring that decisions with social and ethical implications are openly discussed with stakeholders and with civil society (see p. 53).

Additional financial resources

The cost estimates above are based on the financial instruments available in the 7th Framework Programme and their possible evolution under Horizon 2020. However, the HBP-PS has ascertained that national funding agencies, scientific institutions and industrial partners are willing to provide ad-



ditional financial and in-kind contributions if the HBP is approved. Our current best estimate of these commitments is Eur 300 million, spread over the duration of the project – an average of Eur 30 million per year.

The HBP-PS has collected many commitments from potential sources of public funding. In particular, the Swiss federal government plans to support research on brain simulation with funding of Eur 160 million over the next ten years. Forschungszentrum Jülich (Germany) intends to raise Eur 130 million in the context of the centre’s commitment to PRACE, using these funds to build the HBP supercomputing infrastructure in Jülich. Also in Germany, the State of Baden-Württemberg has finalised planning of a new building dedicated to neuromorphic research in the HBP. Many other sources have indicated they would support the HBP, if approved.

The HBP-PS also collected approximately Eur 18 million of commitments from 27 SMEs in 12 European countries and the USA (one company) (see Figures 42 and 43). Several large companies working in medicine, pharmaceuticals, biotechnology and ICT have also expressed strong support for the project, promising to provide matching funding if it is approved. On this basis, we estimate that it would be possible to obtain approximately Eur 50 million of company funding over the duration of the project. Given that this funding would be for specific research projects using HBP facilities, it should be seen as additional to, rather than a replacement for, other sources of funding.

Know-how and technical facilities

In addition to direct grants, many important European and international facilities and large-scale research projects have indicated that they would be prepared to make in-kind contributions to the project, opening their doors for joint research. One of the most significant is the Allen Institute in the USA. Table 5 provides a preliminary picture of some of the support the HBP could expect to receive.

Area	Facility	Contribution to the HBP
Neuroscience		The Allen Institute [4] is coordinating a series of major initiatives to create interactive atlases of the mouse and human brains. The Institute would be a partner in the HBP, providing data of vital importance for the construction of brain models. The HBP in turn would contribute data and tools for Allen Institute brain atlases
		GeneNetwork [179] is a large-scale NIH-funded genetics project that has produced massive genome-to-phenome data sets and analysis tools. The system is used by a large international community with nodes in Holland, Germany, and Switzerland and is already partnering with several EU efforts (COST SysGeNet) and a team at the EPFL. GeneNetwork would make its entire open source system available to the HBP, which would use GeneNetwork tools in its planned mouse cohort study

>>

		The Mouse Brain Library (MBL) [180] is an NIH project that has assembled high-resolution neurohistological data for over 200 mouse genotypes, using them to study the genetic basis for individual differences in brain structure and function. The HBP mouse cohort study would contribute deep neuroimaging data to the MBL and use MBL data for studies of neuronal populations and their relationship to mouse brain function
		The NIH Human Connectome Project [181] is a massive genetic MRI and functional testing study of a cohort of 1200 young adult siblings and identical twins, planned to reach completion in 2015. The HBP would collaborate with the project, integrating human connectome data (including gene sequences) into the HBP GeneNetwork module and into HBP brain models
		NeuroSpin [182] is a large-scale initiative, funded by the French CEA, with the aim of applying the full power of MRI to understanding the nervous system in animals and humans. The NeuroSpin campus, at Saint-Aubin/Saclay near Paris provides researchers with outstanding MRI/MRS equipment and related tools and an advanced computer platform. NeuroSpin houses a 3T and a 7T wide bore MR scanner for clinical studies, as well as a 11.7T wide bore system a 17T small bore system for preclinical studies. In the HBP, NeuroSpin would provide the imaging platform for experiments in cognitive neuroscience
Brain Simulation		The EPFL Blue Brain Project [183] has played a pioneering role in the development of biological detailed brain models. The Brain Simulation Platform would build on the tools and workflows developed in this work
High Performance Computing		The FZJ [184] is a European leader in high performance computing. The centre's Jugene supercomputer is one of the most powerful machines in Europe. FZJ coordinates PRACE [185], a major initiative to coordinate European exascale computing. FZJ would host the main supercomputer used for HBP brain simulation and gradually expand towards exascale capabilities
		The CSCS is the Swiss national supercomputing centre [186]. CSCS would host the machine that HBP researchers would use to develop and test codes for brain simulation
		The Barcelona Supercomputing Center [187] is a major supercomputing centre dedicated to molecular level simulations. In the HBP, the BSC would host molecular dynamics simulations, used to generate data for brain models
		CINECA is the Italian national supercomputing centre and provides Tier 0 hosting for PRACE. CINECA has offered to host the HBP supercomputing facility for massive data analytics
Neuromorphic Computing		Founded in 1994, the Heidelberg ASIC Laboratory for Microelectronics is jointly operated by two university institutes and the Max-Planck-Institute for Nuclear Physics (MPI-K). It offers clean room facilities for chip handling and mounting, advanced test equipment for high speed, low noise measurements as well as a full design, simulation and verification suite for chip design
Neurorobotics		The Department for Automation at the UPM [188] has an array of industrial robots, teleoperation systems, mobile systems, as well as comprehensive facilities and expertise for robotics research
		The eventLab at UB is equipped with a CAVE (Cave Automatic Virtual Environment), a multi-person, high-resolution, virtual reality environment. Also available are full-body motion capture systems, physiological sensing devices and EEG recording systems
		The Robotics and Embedded Systems group at TUM's Department of Informatics is one of the largest and most well equipped robotics departments in Germany. The department specialises in the development of systems with innovative capabilities for perception, cognition, action and control

Table 5: Key technical facilities available to the HBP

Initial HBP Consortium

European scientists and the Human Brain Project

In the early phases of the HBP-PS, we contacted a large number of scientists potentially interested in the Human Brain Project, asking them to propose research tasks for inclusion in the project work plan. The overwhelming response included more than 500 written proposals and an even greater number of informal suggestions. On this basis, the HBP-PS identified 256 leading scientists, representing the full range of disciplines required by the project and coming from many different schools of thought (see Figure 38). This report is the result of their work. All the scientists listed below have agreed to participate in the HBP if it is approved as a FET Flagship Project (see Figure 39).

Molecular and Cellular Neuroscience

- **Grant, Seth, University of Edinburgh (UK)**
- **DeFelipe, Javier, Universidad Politécnica de Madrid (ES)**
- **Araque, Alfonso, Instituto Cajal - Consejo Superior de Investigaciones Científicas (ES)**
- **Armstrong, Douglas, University of Edinburgh (UK)**
- **Canals, Santiago, Instituto de Neurociencias de Alicante UMH-CSIC (ES)**
- **Clasca, Francisco, Universidad Autónoma de Madrid (ES)**
- **De Zeeuw, Chris, Universitair Medisch Centrum Rotterdam (NL)**
- **Freund, Tamas, Magyar Tudományos Akadémia (HU)**
- **Giaume, Christian, Institut national de la santé et de la recherche médicale (FR)**
- **Kleinfeld, David, University of California, San Diego (US)**
- **Koch, Christof, The Allen Institute for Brain Science (US)**
- **Kopanitsa, Maksym, Synome Ltd (UK)**
- **Lujan, Rafael, Universidad de Castilla-La Mancha (ES)**
- **Magistretti, Pierre, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Marin, Oscar, Instituto de Neurociencias de Alicante UMH-CSIC (ES)**
- **Martinez, Salvador, Instituto de Neurociencias de Alicante UMH-CSIC (ES)**
- **Pavone, Francesco, European Laboratory for Non Linear Spectroscopy (IT)**
- **Ponting, Chris, University of Oxford (UK)**
- **Saksida, Lisa, University of Cambridge (UK)**
- **Scattoni, Maria Luisa, Istituto Superiore di Sanità (IT)**
- **Smit, August B., Vrije Universiteit Amsterdam (NL)**
- **Soriano, Eduardo, Institute for Research in Biomedicine Barcelona (ES)**
- **Spijker, Sabine, Vrije Universiteit Amsterdam (NL)**
- **Stampanoni, Marco, Paul Scherrer Institut (CH)**
- **Thomson, Alex, University of London (UK)**
- **Voet, Thierry, Katholieke Universiteit Leuven (BE)**
- **Wang, Yun, Wenzhou Medical College (CN)**
- **Weber, Bruno, Universität Zürich (CH)**
- **Williams, Robert, University of Tennessee (US)**

Cognitive Neuroscience

- **Dehaene, Stanislas, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Amunts, Katrin, Forschungszentrum Jülich (DE)**
- **Axer, Markus, Forschungszentrum Jülich (DE)**
- **Blanke, Olaf, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Born, Jan, Universität zu Lübeck (DE)**
- **Burgess, Neil, University College London (UK)**
- **Chandran, Sidharthan, University of Edinburgh (UK)**
- **Costa, Rui, Champalimaud Foundation (PT)**
- **Dehaene, Ghislaine, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Dudai, Yadin, Weizmann Institute (IL)**
- **Eickhoff, Simon, Universität Düsseldorf (DE)**
- **Frégnac, Yves, Centre national de la recherche scientifique (FR)**
- **Fries, Pascal, Ernst Strüngmann Institute (DE)**
- **Giese, Martin, Universität Tübingen (DE)**
- **Hagoort, Peter, Max Planck Institute for Psycholinguistics (DE)**
- **Hari, Riitta, Aalto-yliopisto (FI)**
- **Herzog, Michael, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Karni, Avi, University of Haifa (IL)**
- **Laurent, Gilles, Max Planck Institute for Brain Research (DE)**
- **Le Bihan, Denis, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Mainen, Zachary, Champalimaud Foundation (PT)**
- **Malach, Rafael, Weizmann Institute (IL)**
- **Mangin, Jean-François, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Nieder, Andreas, Universität Tübingen (DE)**
- **Nyberg, Lars, Umeå Universitet (SE)**
- **Pallier, Christophe, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Parkkonen, Lauri, Aalto-yliopisto (FI)**
- **Paz, Rony, Weizmann Institute (IL)**
- **Pessiglione, Matthias, Institut du cerveau et de la moëlle épinière (FR)**
- **Pinel, Philippe, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Poupon, Cyril, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Robbins, Trevor, University of Cambridge (UK)**
- **Sigman, Mariano, Universidad de Buenos Aires (AR)**
- **Thirion, Bertrand, Institut national de recherche en informatique et en automatique (FR)**
- **Ullmann, Shimon, Weizmann Institute (IL)**
- **van Wassenhove, Virginie, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Yokoyama, Charles, Riken Brain Science Institute (JP)**

Theoretical Neuroscience

- **Destexhe, Alain, Centre national de la recherche scientifique (FR)**
- **Gerstner, Wulfram, Ecole Polytechnique Fédérale de Lausanne (CH)**

- Brunel, Nicolas, Centre national de la recherche scientifique - Université Paris Descartes (FR)
- Deco, Gustavo, Institució Catalana de Recerca i Estudis Avançats (ES)
- Einevoll, Gaute, Universitetet for miljø- og biovitenskap (NO)
- Faugeras, Oliver, Institut national de recherche en informatique et en automatique (FR)
- Hess Bellwald, Kathryn, Ecole Polytechnique Fédérale de Lausanne (CH)
- Jirsa, Viktor, Centre national de la recherche scientifique (FR)
- Maass, Wolfgang, Technische Universität Graz (AT)
- Morrison, Abigail, Universität Freiburg (DE)
- Schrauwen, Benjamin, Universiteit Gent (BE)
- Senn, Walter, Universität Bern (CH)
- Sporns, Olaf, University of Indiana (US)
- Tsodyks, Misha, Weizmann Institute (IL)
- Wierstra, Daan, Deepmind Technologies (UK)
- De Los Rios, Paolo, Ecole Polytechnique Fédérale de Lausanne (CH)
- De Schutter, Erik, Okinawa Institute of Science and Technology (JP)
- Diesmann, Markus, Forschungszentrum Jülich (DE)
- Grubmüller, Helmut, Max-Planck-Institut für Biophysik (DE)
- Hausser, Michael, University College London (UK)
- Hines, Michael, Yale University (US)
- Jérusalem, Antoine, Instituto Madrilenio De Estudios Avanzados (ES)
- Jonas, Peter, Institute of Science and Technology Austria (AT)
- Laio, Alessandro, La Scuola Internazionale Superiore di Studi Avanzati (IT)
- Lavery, Richard, Université de Lyon (FR)
- Lindahl, Erik, Stockholms Universitet (SE)
- Migliore, Michele, Consiglio Nazionale delle Ricerche (IT)
- Noé, Frank, Freie Universität Berlin (DE)
- Orozco, Modesto, Institute for Research in Biomedicine Barcelona (ES)
- Pottmann, Helmut, King Abdullah University of Science and Technology (SA)
- Roth, Arnd, University College London (UK)
- Röthlisberger, Ursula, Ecole Polytechnique Fédérale de Lausanne (CH)
- Sansom, Mark, University of Oxford (UK)
- Schürmann, Felix, Ecole Polytechnique Fédérale de Lausanne (CH)
- Segev, Idan, Hebrew University of Jerusalem (IL)
- Shillcock, Julian, Ecole Polytechnique Fédérale de Lausanne (CH)
- Taly, Antoine, Centre national de la recherche scientifique (FR)
- Tarek, Mounir, National Institute of Standards and Technology Centre for Neutron Research (US)
- Tass, Peter, Forschungszentrum Jülich (DE)
- Triller, Antoine, Institut national de la santé et de la recherche médicale - Ecole normale supérieure (FR)
- Wade, Rebecca, Ruprecht-Karls-Universität Heidelberg (DE)

Neuroinformatics

- **Grillner, Sten, Karolinska Institutet (SE)**
- Baumela, Luis, Universidad Politécnica de Madrid (ES)
- Bjaalie, Jan, Universitetet i Oslo (NO)
- Fua, Pascal, Ecole Polytechnique Fédérale de Lausanne (CH)
- Goebel, Rainer, Maastricht University (NL)
- Grün, Sonja, Forschungszentrum Jülich (DE)
- Hill, Sean, Karolinska Institutet (SE)
- Luthi-Carter, Ruth, University of Leicester (UK)
- Maestu, Fernando, Universidad Complutense de Madrid (ES)
- Martone, Maryann, University of California, San Diego (US)
- Menasalvas, Ernestina, Universidad Politécnica de Madrid (ES)
- Peña, José M., Universidad Politécnica de Madrid (ES)
- Robles, Victor, Universidad Politécnica de Madrid (ES)
- Spiliopoulou, Myra, Universität Magdeburg (DE)
- Tiesinga, Paul, Radboud Uni.Nijmegen (NL)
- van Leeuwen, Cees, Katholieke Universiteit Leuven (BE)
- von Landesberger, Tatiana, Technische Universität Darmstadt (DE)

Brain Simulation

- **Markram, Henry, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Hellgren-Kotaleski, Jeanette, Kungliga Tekniska Högskolan (SE)**
- Andreoni, Wanda, Ecole Polytechnique Fédérale de Lausanne (CH)
- Baaden, Marc, Université Paris Diderot (FR)
- Bernèche, Simon, Swiss Institute of Bioinformatics (CH)
- Carloni, Paolo, German Research School for Simulation Sciences (DE)
- D'Angelo, Egidio, Università degli studi di Pavia (IT)
- Dal Peraro, Matteo, Ecole Polytechnique Fédérale de Lausanne (CH)

Medical Informatics

- **Frackowiak, Richard, Centre Hospitalier Universitaire Vaudois (CH)**
- **Ailamaki, Anastasia, Ecole Polytechnique Fédérale de Lausanne (CH)**
- Ashburner, John, University College London (UK)
- Bogrodzki, Piotr, Politechnika Warszawska (PO)
- Brice, Alexis, Université Pierre et Marie Curie (FR)
- Buneman, Peter, University of Edinburgh (UK)
- Dartigues, Jean-François, Université de Bordeaux (FR)
- Davidson, Susan, University of Pennsylvania (US)
- Draganski, Bogdan, Université de Lausanne (CH)
- Dürr, Alexandra, Université Pierre et Marie Curie (FR)
- Evans, Alan, McGill University (CA)
- Frisoni, Giovanni, Istituto di ricovero e cura a carattere scientifico Fatebenefratelli (IT)
- Gehrke, Johannes, University of Cornell (US)
- Huret, Augustin, EffiScience (FR)

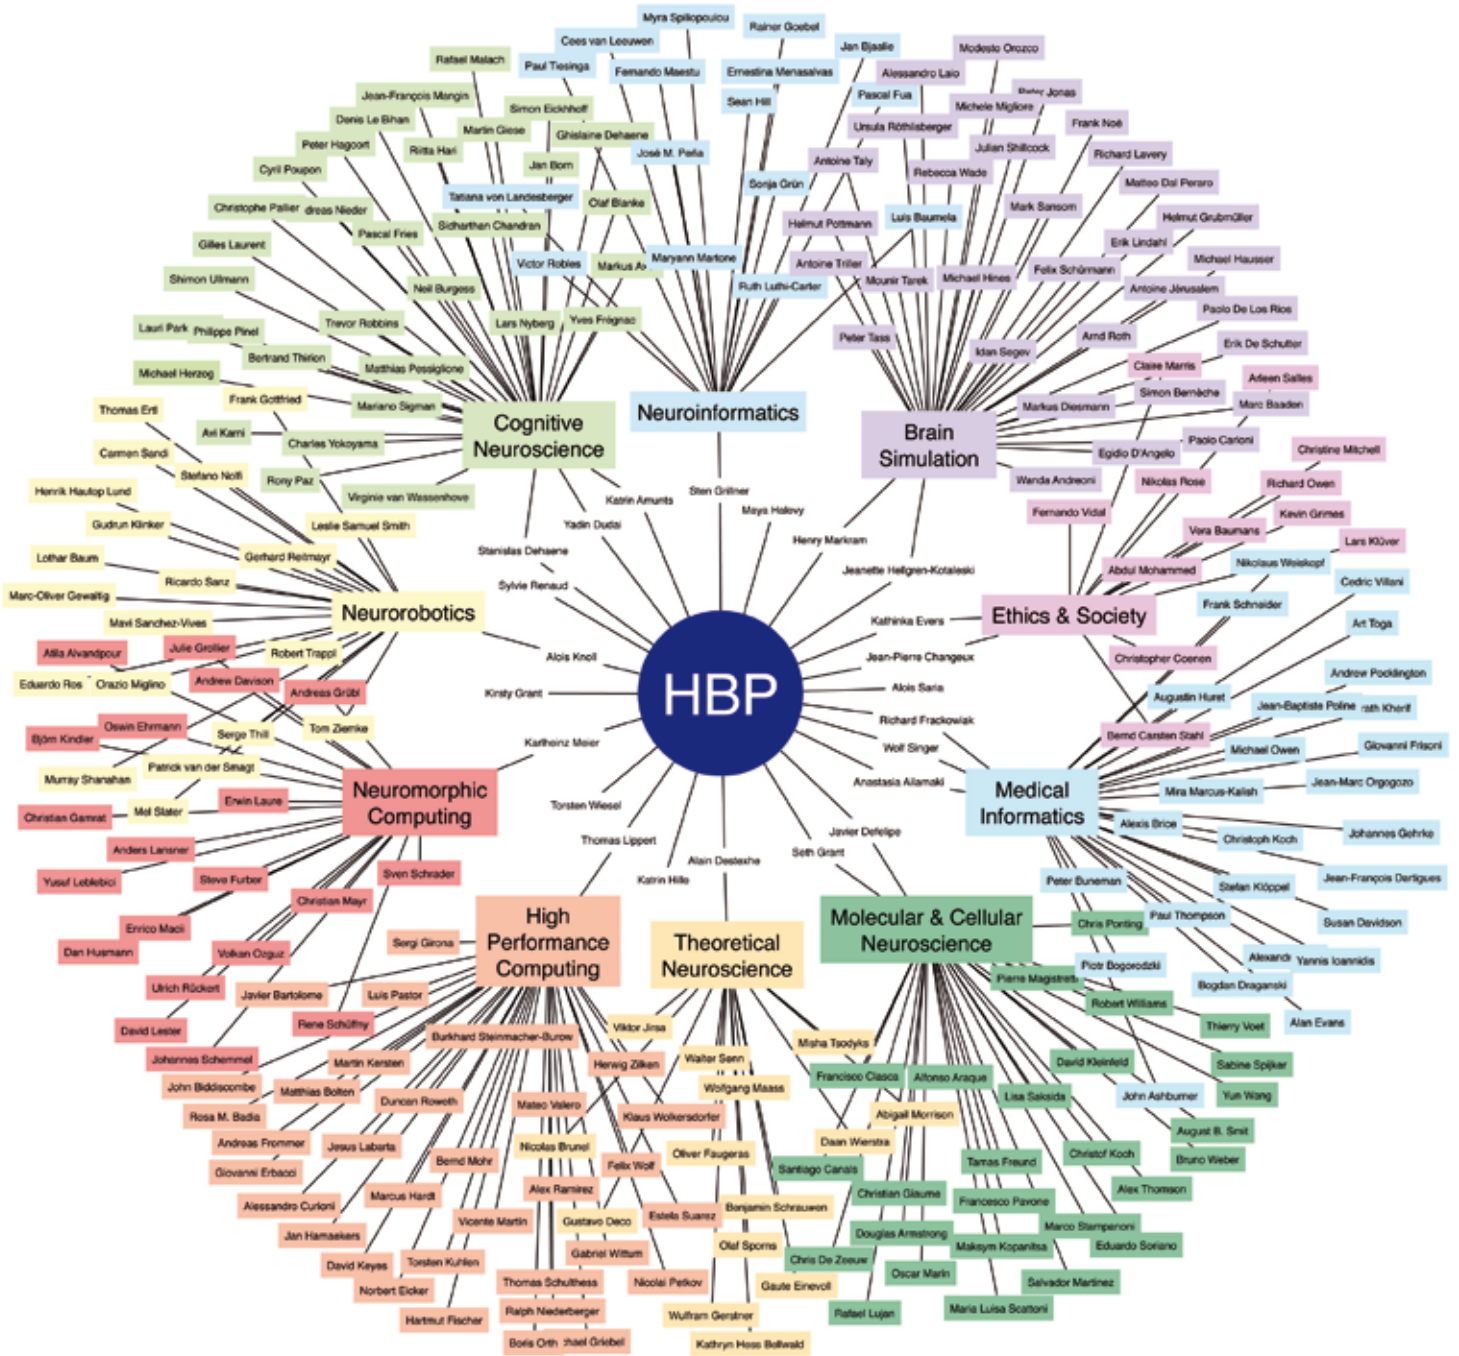


Figure 38: Scientists identified to participate in the HBP. The scientists listed have agreed to participate in the HBP, if it is approved as a FET Flagship

- Ioannidis, Yannis, National and Kapodistrian University of Athens (GR)
- Kherif, Ferath, Centre Hospitalier universitaire Vaudois (CH)
- Klöppel, Stefan, Universitätsklinikum Freiburg (DE)
- Koch, Christoph, Ecole Polytechnique Fédérale de Lausanne (CH)
- Marcus-Kalish, Mira, Tel Aviv University (IL)
- Orgogozo, Jean-Marc, Université de Bordeaux (FR)
- Owen, Michael, Cardiff University (UK)
- Pocklington, Andrew, Cardiff University (UK)
- Poline, Jean-Baptiste, Commissariat à l'énergie atomique et aux énergies alternatives (FR)
- Schneider, Frank, Rheinisch-Westfälische Technische Hochschule Aachen (DE)
- Singer, Wolf, Frankfurt Institute for Advanced Studies (DE)
- Thompson, Paul, University of California, Los Angeles (US)
- Toga, Art, University of California, Los Angeles (US)
- Villani, Cedric, Université de Lyon (FR)
- Weiskopf, Nikolaus, University College London (UK)

High Performance Computing

- **Lippert, Thomas, Forschungszentrum Jülich (DE)**
- *Badia, Rosa M., Barcelona Supercomputing Centre (ES)*
- *Bartolome, Javier, Barcelona Supercomputing Centre (ES)*
- *Biddiscombe, John, Swiss National Supercomputing Centre (CH)*
- *Bolten, Matthias, Bergische Universität Wuppertal (DE)*
- *Curioni, Alessandro, IBM (CH)*
- *Eicker, Norbert, Forschungszentrum Jülich (DE)*
- *Erbacci, Giovanni, Consorzio interuniversitario per la gestione del centro di calcolo elettronico dell'Italia Nord-orientale (IT)*
- *Fischer, Hartmut, Forschungszentrum Jülich (DE)*
- *Frommer, Andreas, Bergische Universität Wuppertal (DE)*
- *Girona, Sergi, Barcelona Supercomputing Centre (ES)*
- *Griebel, Michael, Fraunhofer-Institut für Algorithmen und Wissenschaftliches Rechnen (DE)*
- *Hamaekers, Jan, Fraunhofer-Institut für Algorithmen und Wissenschaftliches Rechnen (DE)*
- *Hardt, Marcus, Karlsruhe Institute of Technology (DE)*
- *Kersten, Martin, Centrum Wiskunde & Informatica (NL)*
- *Keyes, David, King Abdullah University of Science and Technology (SA)*
- *Kuhlen, Torsten, Rheinisch-Westfälische Technische Hochschule Aachen (DE)*
- *Labarta, Jesus, Barcelona Supercomputing Centre (ES)*
- *Martín, Vicente, Universidad Politécnica de Madrid (ES)*
- *Mohr, Bernd, Forschungszentrum Jülich (DE)*
- *Niederberger, Ralph, Forschungszentrum Jülich (DE)*
- *Orth, Boris, Forschungszentrum Jülich (DE)*
- *Pastor, Luis, Universidad Rey Juan Carlos (ES)*
- *Petkov, Nicolai, Universität Groningen (NL)*
- *Ramirez, Alex, Barcelona Supercomputing Centre (ES)*
- *Roweth, Duncan, Cray Inc., The Supercomputer Company (UK)*
- *Schulthess, Thomas, Swiss National Supercomputing Centre (CH)*
- *Steinmacher-Burow, Burkhard, IBM (DE)*
- *Suarez, Estela, Forschungszentrum Jülich (DE)*
- *Valero, Mateo, Barcelona Supercomputing Centre (ES)*
- *Wittum, Gabriel, Goethe Universität (DE)*
- *Wolf, Felix, German Research School for Simulation Sciences (DE)*
- *Wolkersdorfer, Klaus, Forschungszentrum Jülich (DE)*
- *Zilken, Herwig, Forschungszentrum Jülich (DE)*

Neuromorphic Computing

- **Meier, Karlheinz, Ruprecht-Karls-Universität Heidelberg (DE)**
- **Furber, Steve, The University of Manchester (UK)**
- *Alvandpour, Atila, Linköping Universitet (SE)*
- *Davison, Andrew, Centre national de la recherche scientifique (FR)*

- *Ehrmann, Oswin, Fraunhofer-Institut für Zuverlässigkeit und Mikrointegration (DE)*
- *Gamrat, Christian, Commissariat à l'énergie atomique et aux énergies alternatives (FR)*
- *Grollier, Julie, Centre national de la recherche scientifique (FR)*
- *Grübl, Andreas, Ruprecht-Karls-Universität Heidelberg (DE)*
- *Husmann, Dan, Ruprecht-Karls-Universität Heidelberg (DE)*
- *Kindler, Björn, Ruprecht-Karls-Universität Heidelberg (DE)*
- *Lansner, Anders, Kungliga Tekniska högskolan (SE)*
- *Laure, Erwin, Kungliga Tekniska högskolan (SE)*
- *Leblebici, Yusuf, Ecole Polytechnique Fédérale de Lausanne (CH)*
- *Lester, David, The University of Manchester (UK)*
- *Macii, Enrico, Politecnico di Torino (IT)*
- *Mayr, Christian, Technische Universität Dresden (DE)*
- *Ozguz, Volkan, Sabanci Üniversitesi (TK)*
- *Rückert, Ulrich, Universität Bielefeld (DE)*
- *Schemmel, Johannes, Ruprecht-Karls-Universität Heidelberg (DE)*
- *Schrader, Sven, Ruprecht-Karls-Universität Heidelberg (DE)*
- *Schüffny, Rene, Technische Universität Dresden (DE)*

Neurorobotics

- **Knoll, Alois, Technische Universität München (DE)**
- *Baum, Lothar, Robert Bosch GmbH (DE)*
- *Ertl, Thomas, Universität Stuttgart (DE)*
- *Gewaltig, Marc-Oliver, Ecole Polytechnique Fédérale de Lausanne (CH)*
- *Gottfried, Frank, SAP AG (DE)*
- *Hautop Lund, Henrik, Danmarks Tekniske Universitet (DK)*
- *Klinker, Gudrun, Technische Universität München (DE)*
- *Miglino, Orazio, Uni. Naples Federico II (IT)*
- *Nolfi, Stefano, Consiglio Nazionale delle Ricerche (IT)*
- *Reitmayr, Gerhard, Technische Universität Graz (AT)*
- *Ros, Eduardo, Universidad de Granada (ES)*
- *Sanchez-Vives, Mavi, Institució Catalana de Recerca i Estudis Avançats (ES)*
- *Sandi, Carmen, Ecole Polytechnique Fédérale de Lausanne (CH)*
- *Sanz, Ricardo, Universidad Politécnica de Madrid (ES)*
- *Shanahan, Murray, Imperial College London (UK)*
- *Slater, Mel, Universitat de Barcelona (ES)*
- *Smith, Leslie Samuel, University of Stirling (UK)*
- *Thill, Serge, Högskolan i Skövde (SE)*
- *Trapp, Robert, Österreichische Studiengesellschaft für Kybernetik (AT)*
- *van der Smagt, Patrick, Deutschen Zentrums für Luft- und Raumfahrt (DE)*
- *Ziemke, Tom, Högskolan i Skövde (SE)*

Society and Ethics

- **Changeux, Jean-Pierre, Institut Pasteur (FR)**
- **Evers, Kathinka, Uppsala universitet (SE)**
- **Baumans, Vera, Universiteit Utrecht (NL)**
- **Carsten Stahl, Bernd, De Montfort University (UK)**
- **Coenen, Christopher, Karlsruhe Institute of Technology (DE)**
- **Dudai, Yadin, Weizmann Institute (IL)**
- **Grimes, Kevin, Karolinska Institutet (SE)**
- **Klüver, Lars, Danish Board of Technology (DK)**
- **Marris, Claire, Kings College (UK)**
- **Mitchell, Christine, Harvard University (US)**
- **Mohammed, Abdul, Linnéuniversitetet (SE)**
- **Owen, Richard, University of Exeter (UK)**
- **Rose, Nikolas, Kings College (UK)**
- **Salles, Arleen, Centro de Investigaciones Filosóficas (AR)**
- **Vidal, Fernando, Max Planck Institute for the History of Science (DE)**

Management

- **Markram, Henry, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Meier, Karlheinz, Ruprecht-Karls-Universität Heidelberg (DE)**
- **Ailamaki, Anastasia, Ecole Polytechnique Fédérale de Lausanne (CH)**
- **Amunts, Katrin, Forschungszentrum Jülich (DE)**
- **Defelipe, Javier, Universidad Politécnica de Madrid (ES)**
- **Dehaene, Stanislas, Commissariat à l'énergie atomique et aux énergies alternatives (FR)**
- **Destexhe, Alain, Centre national de la recherche scientifique (FR)**
- **Dudai, Yadin, Weizmann Institute (IL)**
- **Evers, Kathinka, Uppsala universitet (SE)**
- **Frackowiak, Richard, Centre Hospitalier universitaire Vaudois (CH)**

Researchers by division

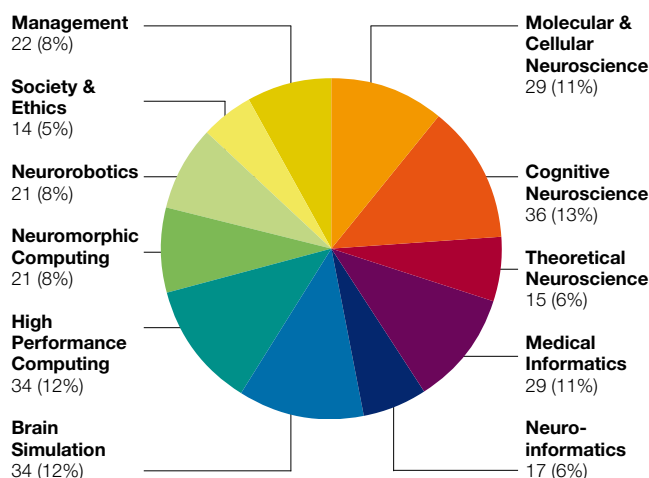


Figure 39: Distribution of researchers by division. Numbers do not include researchers expected to join the project through open calls

- **Grant, Kirsty, Centre national de la recherche scientifique (FR)**
- **Grant, Seth, University of Edinburgh (UK)**
- **Grillner, Sten, Karolinska Institutet (SE)**
- **Halevy, Maya, Bloomfield Science Museum (IL)**
- **Hellgren-Kotaleski, Jeanette, Kungliga Tekniska högskolan (SE)**
- **Hille, Katrin, Universität Ulm (DE)**
- **Knoll, Alois, Technische Universität München (DE)**
- **Lippert, Thomas, Forschungszentrum Jülich (DE)**
- **Renaud, Sylvie, Centre national de la recherche scientifique (FR)**
- **Saria, Alois, Medizinische Universität Innsbruck (AT)**
- **Singer, Wolf, Frankfurt Institute for Advanced Studies (DE)**
- **Wiesel, Torsten, Rockefeller University (US)**

Participating institutions

During the HBP-PS, we identified 150 leading research institutes and universities with the collective know-how to lead the initial phase of the project (see Figure 40). These potential partners cover 24 countries, including all major EU Member States as well as Switzerland, the USA, Japan and China. Table 6 lists these institutions.

Leveraging the strengths and diversity of European research

The HBP must leverage the strengths and diversity of European research. These should be reflected in the composition of the consortium that undertakes the initial work, in the project's governance and management model and in its use of publicly funded financial and technical resources. The consortium that undertakes the initial work should include scientists from different schools of thought with a broad range of theoretical and experimental approaches. The management should ensure that the project has sufficient strategic and administrative flexibility to accommodate new concepts and technologies.

It is especially important that the project provides strong support to the whole scientific community, acting as an incubator for emerging ideas. To this end, we propose that the project dedicate at least 20% of its budget to two programmes explicitly designed to achieve this objective.

The first, modelled on the highly successful FET Programme, should provide support to groups and projects, selected through open calls for proposals and independent peer review; the second should follow the example of the ERC and Marie-Curie programmes, providing Advanced Research Grants, Young Investigator Grants and Studentships to individual researchers at different stages of their careers (see Figure 41).

Support for individual researchers

This programme, based on an open competitive process, would provide awards supporting individual researchers at different stages in their careers.

- Approximately ninety one-year *Advanced Research Grants* (8-10 per year) would go to internationally recognised

Institution	Country	Institution	Country	Institution	Country
Centro de Investigaciones Filosoficas	AR	Universität Tübingen	DE	Università degli Studi di Napoli	IT
Universidad de Buenos Aires	AR	Universitätsklinikum Freiburg	DE	Federico II	
Innsbruck Medical University	AT	Bergische Universität Wuppertal	DE	Università degli studi di Pavia	IT
Institute of Science and Technology Austria	AT	Rheinisch-Westfälische Technische Hochschule Aachen	DE	Okinawa Institute of Science and Technology	JP
Österreichische Studiengesellschaft für Kybernetik	AT	Teknologirådet	DK	Riken Brain Science Institute	JP
Technische Universität Graz	AT	Danmarks Tekniske Universitet	DK	Centrum Wiskunde & Informatica	NL
Katholieke Universiteit Leuven	BE	Barcelona Supercomputing Centre	ES	Universitair Medisch Centrum Rotterdam	NL
Universiteit Gent	BE	Universidad Complutense de Madrid	ES	Vrije Universiteit Amsterdam	NL
McGill University	CA	Instituto Cajal - Consejo Superior de Investigaciones Científicas	ES	Universit�t Groningen	NL
Centre Hospitalier universitaire Vaudois	CH	Instituci� Catalana de Recerca i Estudis Avan�ats	ES	Radboud Universiteit Nijmegen	NL
Swiss National Supercomputing Centre	CH	IMDEA Materials institute (Madrid Institute for Advanced Studies of Materials)	ES	Maastricht University	NL
Ecole Polytechnique F�d�rale de Lausanne	CH	Institute for Research in Biomedicine Barcelona	ES	Universiteit Utrecht	NL
IBM	CH	Instituto de Neurociencias de Alicante	ES	Universitetet for milj�- og biovitenskap	NO
Paul Scherrer Institut	CH	UMH-CSIC	ES	Universitetet i Oslo	NO
Swiss Institute of Bioinformatics	CH	Universidad Aut�noma de Madrid	ES	Politechnika Warszawska	PO
Universit�t Bern	CH	Universidad de Castilla-La Mancha	ES	Champalmaud Foundation	PT
Universit�t Z�rich	CH	Universidad de Barcelona	ES	King Abdullah University of Science and Technology	SA
Universit� de Lausanne	CH	Universidad de Granada	ES	Karolinska Institutet	SE
Wenzhou Medical College	CN	Universidad Polit�cnica de Madrid	ES	Kungliga Tekniska h�gskolan	SE
Universit�tsklinikum Aachen	DE	Universidad Rey Juan Carlos	ES	Link�ping Universitet	SE
Universit�t Bielefeld	DE	Aalto-yliopisto	FI	Linn�universitetet	SE
Robert Bosch gmbH	DE	Commissariat � l'�nergie atomique et aux �nergies alternatives	FR	Stockholms Universitet	SE
Deutsches Zentrum f�r Luft- und Raumfahrt	DE	Centre national de la recherche scientifique	FR	H�gskolan i Sk�vde	SE
Technische Universit�t Dresden	DE	EffiScience	FR	Ume� Universitet	SE
Ernst Str�ngmann Institute	DE	�cole normale sup�rieure	FR	Uppsala universitet	SE
Frankfurt Institute for Advanced Studies	DE	Institut du cerveau et de la moelle �pini�re	FR	Sabancı �niversitesi	TK
Goethe Universit�t	DE	Institut national de recherche en informatique et en automatique	FR	Cardiff University	UK
Fraunhofer-Institut f�r Zuverl�ssigkeit und Mikrointegration	DE	Institut national de la sant� et de la recherche m�dicale	FR	Crax Inc.	UK
Fraunhofer-Institut f�r Algorithmen und Wissenschaftliches Rechnen	DE	Institut Pasteur	FR	De Montfort University	UK
Freie Universit�t Berlin	DE	Universit� de Bordeaux	FR	Deepmind Technologies	UK
Forschungszentrum J�lich	DE	Universit� de Lyon	FR	Kings College	UK
German Research School for Simulation Sciences	DE	Universit� Paris Diderot	FR	Imperial College London	UK
Heidelberg Institute for Theoretical Studies	DE	Universit� Pierre et Marie Curie	FR	The University of Manchester	UK
Ruprecht-Karls-Universit�t Heidelberg	DE	National and Kapodistrian University of Athens	GR	Synome Ltd	UK
Karlsruhe Institute of Technology	DE	Magyar Tudom�nyos Akad�mia	HU	University College London	UK
Max-Planck-Institut f�r Biophysik	DE	Hebrew University of Jerusalem	IL	University of Cambridge	UK
Max Planck Institute for Brain Research	DE	Bloomfield Science Museum Jerusalem	IL	University of Edinburgh	UK
Max Planck Institute for Psycholinguistics	DE	Tel Aviv University	IL	University of Exeter	UK
Max Planck Institute for the History of Science	DE	University of Haifa	IL	University of Leicester	UK
SAP AG	DE	Weizmann Institute	IL	University of London	UK
Technische Universit�t Darmstadt	DE	Consorzio interuniversitario - CINECA	IT	University of Oxford	UK
Technische Universit�t M�nchen	DE	Consiglio Nazionale delle Ricerche	IT	University of Stirling	UK
Universit�t Ulm	DE	Istituto Superiore di Sanit�	IT	University of Oxford	UK
Universit�t D�sseldorf	DE	Istituto di ricovero e cura a carattere scientifico Fatebenefratelli	IT	The Allen Institute for Brain Science	US
Universit�t Freiburg	DE	European Laboratory for Non Linear Spectroscopy	IT	Cornell University	US
Universit�t zu L�beck	DE	Politecnico di Torino	IT	Harvard University	US
Universit�t Magdeburg	DE	La Scuola Internazionale Superiore di Studi Avanzati	IT	University of Indiana	US
Universit�t Stuttgart	DE			National Institute of Standards and Technology Center for Neutron Research	US
				Rockefeller University	US
				University of California, Los Angeles	US
				University of California, San Diego	US
				University of Pennsylvania	US
				University of Tennessee	US
				Yale University	US

Table 6: Institutions participating in the initial HBP Consortium

Initial HBP Consortium

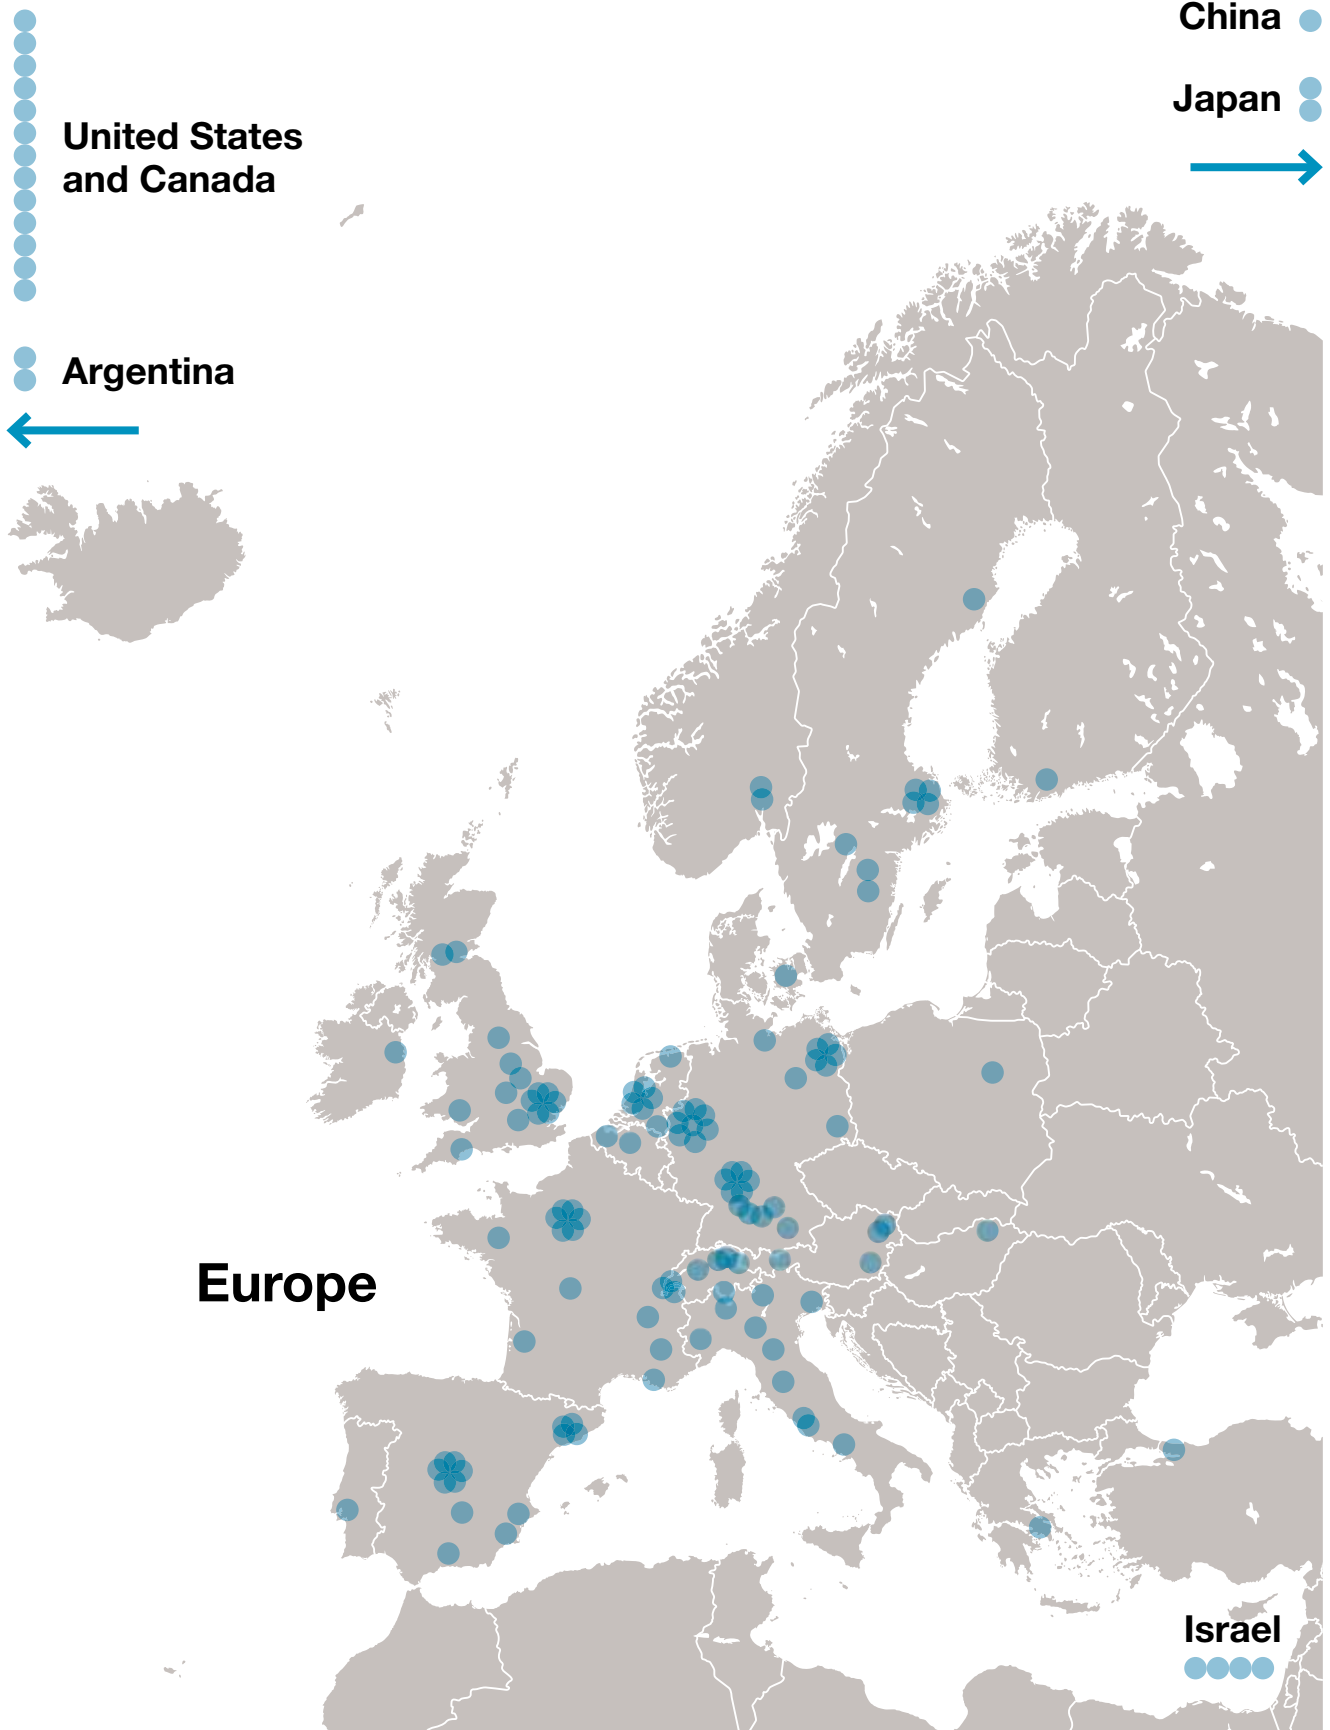


Figure 40: Research institutions identified to participate in the HBP, if approved as a FET Flagship Project

scholars wishing to perform independent interdisciplinary research using the HBP platforms. The value of the awards (Eur 500.000/year) would cover the salary of the researcher, salaries for two PhD students, travel money and lab equipment. The cost would amount to approximately 20% of the total funding for individual researchers.

- About sixty five-year awards would go to *Young Investigators* (5-7 per year): advanced post-doctoral researchers wishing to pursue independent research projects in labs managed by one of the HBP partners. The awards (Eur 300.000/year) would cover the salary for the recipient, for two PhD students supporting his/her work, as well as for research equipment and consumables. The cost would amount to approximately 40% of the funding available.
- One hundred and ninety three-year *Post-doctoral fellowships* (15-20 per year) of Eur 100.000/year would be assigned to post-doctoral researchers wanting to enter independent research for the first time. The cost would amount to approximately 25% of the total funding.
- Two hundred and ninety three-year *HBP studentships* (220-230 per year) of Eur 50.000/year would be reserved for PhD students. The cost would amount to approximately 15% of the funding.

Support for groups and projects

This programme would be funded by the planned ERANET+ or ERANET and would provide support to projects proposed and managed by research groups from outside the initial Consortium. Projects selected under these schemes should receive between Eur 100.000 and Eur 300.000 of funding per year for up to three years. Projects would be selected via competitive calls for proposals addressing all four areas of HBP research (data, theory, platforms and applications). Evaluation would be by independent peer review.

Gender equality

Research in the Human Brain Project would include work in a broad range of disciplines from molecular and cellular biology to mathematics, medicine, computer engineering, robotics and even philosophy. These disciplines are characterised by widely differing rates of female participation, often with large variations between countries. Nonetheless two tendencies stand out. First, in all major disciplinary areas, except engineering, manufacturing and construction, at least 40% of new PhDs in the EU-27 countries are female [189] (data for 2006). Second, in these same countries, the proportion of women decreases with each successive level from students to researchers to professors [190]. In 2007, only 19% of senior Grade A positions were occupied by women [189]. This tendency is confirmed by the HBP-PS: of the Principal Investigators who expressed interest in participating in the HBP, only 15% were women. A recent EU report suggests that the cause of this inequality is “largely on the demand side, that is derived from employer policies and/or strategies” [190]. In view of these findings, and of

Promoting different schools of thought

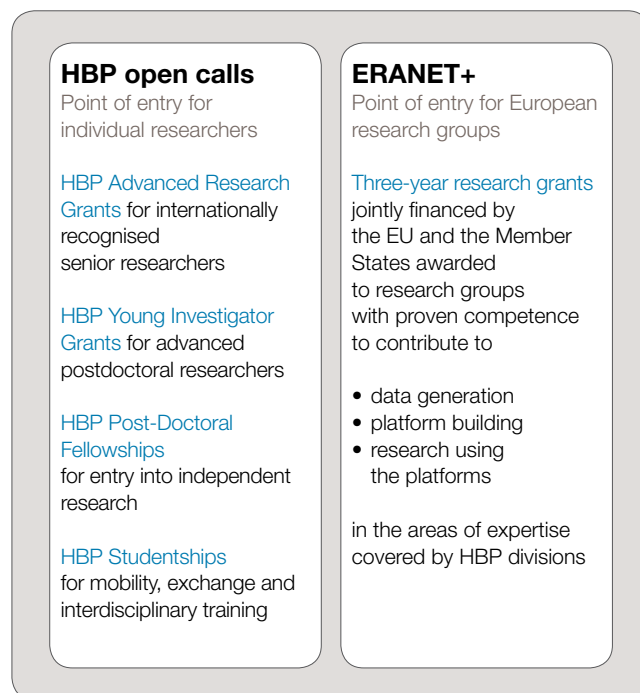


Figure 41: Support for individual researchers and research groups in the HBP

the high profile of flagships in European research, it is important that the HBP should play a pioneering role in the recruitment of female researchers, in particular to Principal Investigator positions.

We therefore recommend that the project define annual targets for the proportion of female researchers at different levels within the project (PhD students, post docs, task and work package leaders, division leaders, senior management positions) and for different project activities (core project research, research projects supported under open calls for proposals, research grants and studentships, management).

Targets should ensure a steady year on year increase in the proportion of female researchers working in the project. In the case of senior researchers, the goal should be to increase the proportion of female Principal Investigators by between 2% and 3% every year. By the end of the project between 30% and 40% of HBP Principal Investigators should be women.

Alignment with national, European and international priorities

High-level priorities

Despite differences of emphasis and language, European governments and the European Union share broadly similar research priorities. Although not all governments explicitly declare their priorities, all place great emphasis on three themes: health; energy and the environment; and improving competitiveness and employment.

HBP: A unique shared R&D platform for industry



Figure 42: Companies that have expressed an interest in participating in the HBP

In each of these areas, we have analysed the national research priorities defined in political agreements among the political parties (e.g. Germany: [191]); National Research Programmes (e.g. France:[192] Italy: [193], Spain: [194]); reports by national research councils (e.g. UK: [195]), statements by policy-makers (Austria, Denmark, France, Israel, Switzerland); and the implicit priorities expressed in patterns of research investment and output (Germany, Switzerland, the UK). We have also analysed the emerging consensus among policy-makers about the need for stronger investment in brain research. From our analysis, we conclude that the HBP is well aligned to current and emerging priorities.

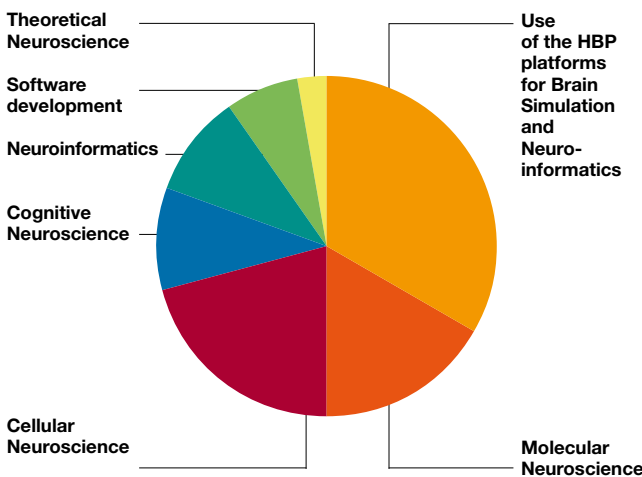
Health research

In all the countries we examined, policy-makers assign a high priority to medical research. A frequently recurring

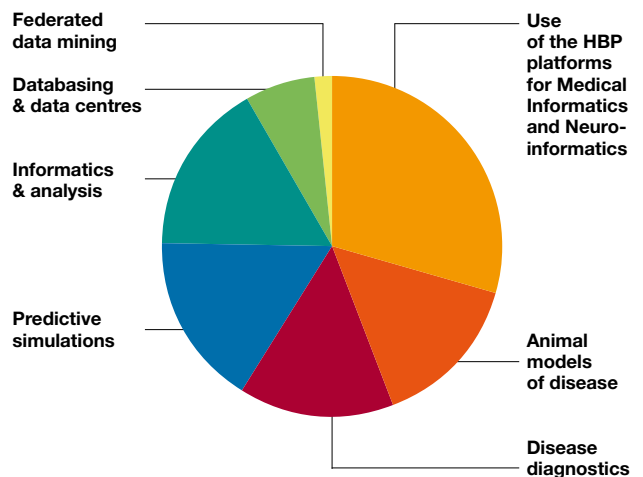
theme is the need to respond to the aging of the European population (the British Medical Research Council lists dementia as top priority for research [195]). Two other important areas of interest are the relationship between genetics and disease [195] and personalised medicine, cited explicitly in the agreement between the political parties which formed the current German government [191].

The HBP speaks to these concerns. HBP work in medical informatics and brain simulation would create a vital new tool for medical researchers, enabling them to federate and mine very large volumes of clinical data, identifying biological signatures of disease, and improving our ability to understand, diagnose and treat diseases of the brain. Some of the earliest results are likely to be in diseases of old age such as Alzheimer's. The HBP would also make an important contribution to basic medical research, proposing completely new strategies for understanding and classifying brain disease and

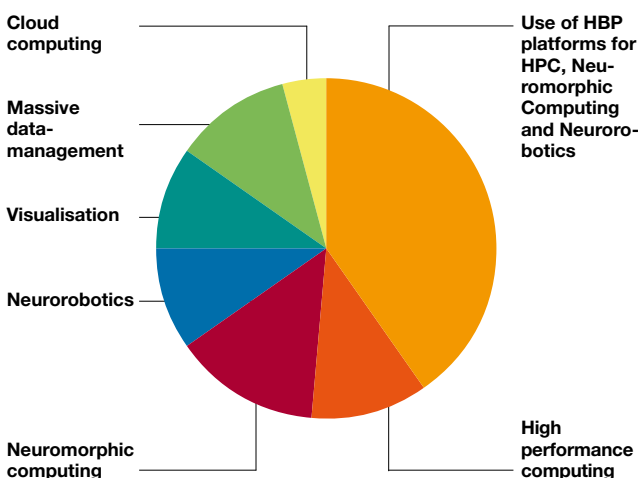
Brain Simulation



Medical Informatics



Future Computing



General Technologies




Figure 43: Which areas of HBP would be most useful for your company's research? Poll responses by industry sector

for personalised medicine. Finally, the proposed HBP mouse cohort study would provide valuable insights into the general principles through which genetic defects affect different levels of brain organisation. In summary, the HBP would make a very significant contribution to health research.

Energy and the environment

In all European countries, research policy places a high priority on energy and the environment. One important element in these policies is the need to make more efficient use of energy, a requirement mentioned both in Italian and in German policy documents [191, 193]. None of these documents explicitly mentions the rapid rise in energy consumption for computing and telecommunications. It is evident, however, that research in this area would closely match national priorities.

As pointed out earlier, the human brain is orders of magnitude more energy-efficient than even the best designed digital computer. An important goal of the HBP would thus be to build very compact, energy-efficient computing systems inspired by the architecture of the brain – and by the way the brain processes and represents information. Such systems would not replace current technology. They could however, become equally important, increasing the services available to industry and citizens while avoiding a parallel, unsustainable increase in energy consumption.

Improving competitiveness and employment

A strategic goal for the Europe 2020 strategy [196] and for all European governments is to use innovation as a lever to improve the competitive position of the European economy.

Several policy documents refer explicitly to particular sectors of research with strong economic potential, many of which coincide with areas in which the HBP is focusing its research. France for example places a high priority on ICT and particularly on high performance computing [192], and has officially confirmed its interest in HBP work in this area. Germany, explicitly mentions microelectronics (at the centre of the HBP’s efforts in neuromorphic computing) and long-term investment in pharmaceuticals (an area in which HBP work in brain simulation and medical informatics can make an important contribution) [191]. Spain [194] and Italy [193] also mention pharmaceuticals. We have received an official statement from the Danish Ministry of Research expressing their interest in HBP work in neurorobotics.

An analysis of actual research investment provides further insights. For example, Germany, Switzerland, Spain and Italy and several other European countries are all making significant investments in supercomputing infrastructure and in new supercomputing technology (including infrastructure explicitly dedicated to brain simulation). Once again the objectives of the HBP are a close match for national priorities.

The brain

Although brain diseases far outweigh even cardiovascular disease and cancer as a burden on European society [1], the brain still has a relatively low priority in European research. There are exceptions. Several national research projects mention the need for more research on dementia. Since 2005, the Swiss federal government has considered EPFL’s Blue Brain Project as a national research priority and has committed to providing long-term support for the project.

Public spending on brain research




Figure 44: Public spending on brain research by country (2005) [1]

Industrial spending on brain research




Figure 45: Industrial spending on brain research by country (2005) [1]

The HBP-PS has received an official statement from the Israeli National Council for R&D that Israel has a strong interest in research into the brain and its diseases [197]. Up to now, however, governments and industry have failed to make the investment in brain research that many believe is essential. In 2005 for instance total spending for research into brain disorders was roughly equivalent to funding for cancer research (brain disorders: Eur 4.1 billion; cancer: Eur 3.9 billion), even though the estimated cost of brain disease was more than three times higher [198].

Today, however, there is a growing consensus that the brain should become an important priority for research. In Europe, for instance, the European Brain Council study of the

economic cost of brain disease [1] has led to numerous calls for increased investment. In the United States, private philanthropists are investing heavily in basic neuroscience [4, 199]. More generally, there is agreement that understanding the brain is vitally important and that we urgently need to overcome the fragmentation that characterises current research [24]. This is precisely the goal of the HBP.

Performance metrics

Managing a project as large and complex as the HBP requires well-defined measures of performance. We propose the following performance metrics.

Category of research	Goal	Metrics/means of proof
All HBP research and development	High scientific quality	Papers accepted in high impact journals and conferences Citations of papers accepted in high impact journals and conferences
	Respect for milestones	Mean and median delay in achieving milestones
	Efficiency of financial planning	$\frac{\text{Expend}_{actual} - \text{Expend}_{planned}}{\text{Expend}_{planned}}$
	Contribution to strategic goals	Publications of results indicated as milestones in the HBP work plan (e.g. first publication of a specific class of brain model, first prototype of a specific class of medical or ICT application)
	Use of HBP work by groups outside the HBP Consortium	Use of results cited in third-party scientific publications
	Media impact	References to the HBP by high impact non-academic media (press, radio, television, important online resources)
Research funded by open calls	Openness to international scientific community	Number of calls Budget allocated to calls Number of projects presented via open calls Number of projects accepted Funding allocated to accepted projects Funding distributed to accepted projects Number of applicants for visitors programme/fellowship programmes % of accepted applicants % of funding to PIs from outside the HBP Consortium
	Support for diversity	% of funding to PIs from smaller countries % of funding to female PIs % of funding to young PIs % of funding for research topics not explicitly mentioned in HBP work plan
	Scientific quality	Papers published by accepted projects (only papers directly related to HBP research, only papers in high impact journals and conferences)
	Contribution to strategic goals of HBP	Use of results cited in papers Project reporting EU project reviews

>>

Category of research	Goal	Metrics/means of proof
	Gender equality	% of female PIs (absolute number, number as % of target) % of female researchers (absolute number, number as % of target) % of female PIs for third-party research projects (absolute number, number as % of target) % of female researchers accepted in young investigators awards (absolute number, number as % of target)
Building and managing the platforms	Compliance with technical specifications	Opinion of reviewers in EU project reviews
	Quality of Service and user satisfaction	Quantitative QoS Indicators (availability etc.), user questionnaires, user focus groups, number of regular users
Using the platforms	Use of the platforms as a community resource	User evaluations Use of platform by research groups in HBP Consortium (numbers of users, hours of use) Use of platform by research groups outside the HBP Consortium (numbers of users, hours of use) Number of companies using platforms (collaborative projects) Number of companies using platforms (paid use of platforms) Revenue from companies using platforms
Society and Ethics	Public participation	Number of participants in HBP ethics events/consultations etc. Number of citations of events/consultations in non-academic media Participant evaluation of events (questionnaires, focus groups)
	Internal ethical awareness	Number of internal ethical workshops Number of participants in internal ethical workshops % of researchers who have participated in an internal ethical workshop in previous 12 months Anonymous participant evaluation of internal ethical workshops
Management	Maintain and encourage interaction between research groups	Number of major project meetings Number of inter-group meetings Number of institutionalised collaborations between groups Participant evaluation (questionnaires, focus groups)
	Quality of administration	Participant evaluation (questionnaires, focus groups) Number of unresolved administrative issues with partners/applicants as % of total issues Number of unresolved administrative issues with EU Commission as % of total issues
	Efficiency of administration	Mean time to prepare calls for proposals Mean time to evaluate proposals Mean time to award research contracts Mean time to process project financial statements Mean time to make payments to selected projects Mean time to distribute EU payments to partners Mean time to complete EU administrative and financial reports Mean time to process requests for information from partners/applicants etc.

Table 7: The HBP: performance metrics

Use of results and dissemination of knowledge

Dissemination

Organisation

The HBP should make a major effort to build awareness of the project's goals and achievements, in relevant scientific communities, as well as among decision-makers and the general public. Dissemination activities should be coordinated by a professionally staffed *Communications Department*, which would collaborate closely with its counterparts in partner institutions.

Dissemination for decision-makers and the general public

The World Wide Web and online media

Online and social media are in a state of continuous evolution making it impossible to plan their use more than a few years in advance.

During preparations for the launch of the project, the HBP Communications Department should create a dedicated web portal. A permanent staff of science journalists should guarantee that the portal is continuously updated with news about the project and about relevant work by other researchers. The site, which should include special sections for children and young adults, would make intensive use of interactive multimedia and provide access to a broad range of educational and tutorial materials, suitable for use in schools and universities.

An important goal should be to encourage discussion and debate about the ethical and social implications of HBP research. Other channels to reach scientific audiences should include science blogs, Facebook pages, as well as live streaming and videos of events and lectures. Plans for the use of these and other online media should be regularly updated as technology evolves.

The general and the specialised media

Relations with the general and specialised media should be managed by the HBP Communications Department, which would build permanent relationships with journalists working in radio, TV, newspapers and magazines, online news sources popular science magazines, online sources of science news etc. The department should ensure that the media is

informed of major project developments and events and coordinate support (video materials, lab visits, interviews etc.) for authors working on documentaries, TV programmes, articles about the project etc.

Science museums

We recommend that the HBP organise a systematic programme of collaboration with European science museums and science fairs, which provide an extremely effective channel to reach the general public, and in particular school children. We suggest that collaboration take the form of permanent, continuously updated exhibits providing current news on the work of the project, as well as larger travelling exhibitions on the brain, brain disease, and brain-inspired technologies (see Figure 46).

Education and ethics programmes

The HBP's education and ethics programmes should play a major role in building awareness of the project's work and results, delivering courses to large numbers of students and involving external ethicists, social scientists and the general public in formal and informal debates on the project's ethical and social implications. In management terms, these programmes would be kept separate from dissemination activities, ensuring they maintain a neutral and, where necessary, a critical stance.

Dissemination for the scientific community

Access to ICT platforms

The ICT platforms should be designed and operated to provide services to researchers, clinicians and technology developers from outside the project. Institutions wishing to guarantee access to a platform for their researchers would be able to do so by agreeing to cover a proportion of the cost.

Access to data and software

Neuroscience data generated by the project should be deposited in brain atlases freely accessible to the scientific community. Software for neuroinformatics should be released in open source format with a license allowing free use for academic research.

Access to neuromorphic hardware

We recommend that the HBP should create low-cost neuromorphic computing systems, making them available to researchers and developers without payment or for a nominal fee.



Museums and outreach




Figure 46: Permanent HBP displays at science museums around the world

The project should leverage community talents and enthusiasm, funding awards and competitions for novel applications.

Scientific publications

The HBP should publish its methods and results in well-established international journals and at leading international conferences. As far as possible, papers should be published in Open Access Journals and/or deposited on pre-print servers. In addition to publications in journals, the project should fund the publishing of a series of monographs dedicated to different aspects of the project. These should include basic neuroscience, brain simulation, medical informatics, neuromorphic computing, neurotechnologies, neurorobotics and ethics.

Conferences

The HBP should organise a series of annual conferences dedicated to specific themes relevant to the project and each including a significant proportion of speakers from outside the HBP.

The World Wide Web and other online media

The project should make intensive use of the World Wide Web, creating and maintaining a high quality web portal, and providing specific sections for scientists and technologists in specific disciplines. Other channels to reach scientific audiences should include science blogs, Facebook pages, as well as live streaming and videos of events and lectures. Plans for the use of new media should be regularly updated as technology evolves.

Exploitation

Research outputs

HBP research outputs with exploitation potential would include:

- Experimental data and research tools; e.g. the 3D atlases and related tools produced by HBP work in neuroinformatics, the brain models, the *Brain Builder* and the *Brain Simulator* developed by the *Brain Simulation Platform*.
- Technical capabilities (e.g. capabilities for the federated mining of distributed sources of clinical data, capabilities for brain, disease and drug simulation).
- Information Technologies (e.g. digital, analogue and hybrid neuromorphic computing systems and devices; components and systems for neurorobotics).
- Prototype applications of these technologies.

Technology transfer

HBP technologies and applications with commercial potential should be used to provide commercial services. Examples include diagnostic services for hospitals, clinical data mining, disease simulation and drug simulation services for pharmaceutical companies and biotech developers. Income from these services would support the project's financial sustainability. Potentially valuable technologies and applications should be patented and licensed to interested third parties for industrial and commercial development.

Management of intellectual property

A major challenge for the Human Brain Project would be the widely dispersed nature of the intellectual property developed by the partners. To handle these issues, we recommend that the HBP should form a licensing company offering one-stop shopping to third parties wishing to license project-generated IP. Universities and other project partners that generate IP would be the owners of the IP, subject to the provisions of relevant national law. Jointly produced IP would be jointly owned. Funds from licensing would be used as a source of funding for the project, helping to guarantee its long-term sustainability.

Risks and contingencies

Category	Risk	Probability	Impact	Evaluation and contingency planning
Political	Political differences between key countries participating in the project	Unknown	High	Conflicts compromising the ability of core partners to participate would probably lead to the halting of the project
	Unpredictable changes in European/national research policy	Unknown	High	Changes in policy compromising the ability of core partners to participate would probably lead to the halting of the project
	Restrictions on travel to Europe for non-European participants	Moderate	Low	Restrictions of this kind would reduce operational efficiency but would not compromise the success of the project
Financial	Inability of Commission/national funding agencies to provide expected level of funding	Moderate	High	The HBP requires very large investment in research and infrastructure. Lack of adequate funding would endanger the community nature of the project and cause major delay. The project should be aware that certain countries would not be able to participate in funding and should make appropriate arrangements to allow participation by researchers in these countries
	Exchange rate fluctuations or price inflation	High	Moderate	The financial estimates in this report are based on fixed 2012 exchange rates. Professional financial management can guarantee some risk protection. However resilience to major currency shocks depends on EU willingness to renegotiate contracts
	Lack of continuity in national and/or Community funding	Moderate	High	Any break in continuity of funding would force partners to reduce personnel, losing vital investment in human capital. This would inevitably cause delays and could in severe cases compromise the viability of the project
	Unforeseen cost overruns in major infrastructure (CAPEX or operating costs)	High	High	This is a major risk in all large-scale projects. In the case of the HBP, major infrastructure investments would be largely managed by national funding agencies, which would use their usual methods to monitor and manage cost overruns
Managerial	Procedural delays due to EU and/or national administrative/financial regulations	Moderate	Low	Experience shows that administrative delays have a major impact on administrations themselves but do not usually affect the overall dynamics of a project
	Conflict within management team and/or among partners	Moderate	Low	Experience with large projects shows that conflict can usually be kept within acceptable levels. In the case of severe conflict, the HBP management structure is sufficiently flexible to take the necessary decisions
Infrastructure	Failure/delay in deploying <i>High Performance Computing Platform</i>	Low	High	The project is essentially dependent on the <i>High Performance Computing Platform</i> . In the early phases of the project partners could use available national facilities. However these facilities do not offer the power to support the later phases of the project

Infrastructure	Failure/delay in deploying <i>Brain Simulation Platform</i>	Low	High	The project is essentially dependent on the availability of the <i>Brain Simulation Platform</i> . Any delay would have a major impact on all project activities
	Failure/delay in deploying <i>Medical Informatics Platform</i>	Low	Moderate	The deployment of the federated data infrastructure depends on decisions by hospitals, which do not belong to the project. Plans for the development of the infrastructure are thus extremely conservative, allowing a long period of time to recruit participating institutions. Failure to deploy the full infrastructure would reduce the scientific value of the data generated. However much useful research could still be done and the failure would not compromise other components of the project
	Failure/delay in deploying <i>Neuromorphic Computing Platform</i>	Low	High	Failure to deploy the <i>Neuromorphic Computing Platform</i> would have a severe impact on all technological development activities and would severely limit possibilities for conducting closed-loop experiments
	Failure/delay in deploying <i>Neurorobotics Platform</i>	Low	High	Failure to deploy the <i>Neurorobotics Platform</i> would make it impossible to conduct closed-loop experiments, compromising a key goal of the project
Scientific/ technological	Failure/delay in generating molecular or cellular data	Low	Moderate	Failure/delay in the generation of molecular/cellular data would reduce the quality of the brain models produced by the <i>Brain Simulation Platform</i> . However there would be no major impact on the development of the technology needed to integrate the data when it becomes available.
	Failure/delay in establishing predictive informatics capabilities	Moderate	Moderate	Failure/delay in establishing planned capabilities in predictive neuroinformatics would reduce the quality of the brain models produced by the <i>Brain Simulation Platform</i> . However there would be no major impact on the development of the technology needed to integrate the data when it becomes available
	Failures/delays in establishing simulation models	Moderate	Moderate	Failure/delay in establishing simulation models would have a major impact on all project activities using the models but would not prevent progress on technology development
	Failures/delays in establishing closed-loop technologies	Moderate	High	Establishing closed-loop technologies is one of the greatest scientific challenges facing the HBP and the risk of failure is real. Failure to establish the technology would have a significant impact on basic research in neuroscience and the development of neurorobotics applications

Table 8: The Human Brain Project: risks and contingency planning

5 Impact

Science

Data and platforms

The fragmentation of neuroscience data and research is widely recognised as the greatest obstacle to achieving an integrated understanding of the relationship between brain structure at different levels of organisation, and behaviour. The data generated by the HBP, in combination with the project's ICT platforms, would bring this goal within reach.

HBP 200 Mouse Cohort. The HBP 200 Mouse Cohort would allow researchers to correlate genetic variations with variations in brain structures at different levels of biological organisation and with variations in behaviour. A parallel human cohort would allow them to compare the effects of genetic variations in mice against similar variations in humans.

HBP data on cognitive neuroscience would describe cognitive tasks, amenable to simulation, and characterise their underlying neuronal architectures, providing guidelines for modelling.

The HBP Neuroinformatics Platform would provide researchers with the advanced tools they need to analyse experimental data on brain structure and brain function, while relieving them of the need to develop and maintain their own software for image processing.

HBP atlases for the mouse and the human brain would make it easier for them to interpret their experimental data in the context of data contributed by other groups (e.g. data for other levels of biological organisation, other brain areas and other species; data for animals of different ages).

HBP tools for predictive neuroinformatics would allow them to mine large volumes of data from different sources, discovering statistical regularities between data for different levels of biological organisation, different brain areas, different species etc. and exploiting them to predict values for data points where experimental measurements are impossible or unavailable. The use of these techniques would maximise the information extracted from lengthy and expensive experiments, avoiding duplication and enabling scientists to fill in their picture of brain structure more rapidly than would otherwise be possible.

Combined with the computing, analysis and visualisation capabilities of the *High Performance Computing Platform*, the *Brain Simulation Platform* would allow researchers to build, simulate and validate models of animal and human brains at every possible scale (abstract computational models, point neuron models, detailed cellular level models of neuronal circuitry, molecular level models of small areas of the brain, multi-scale models that switch dynamically between different levels of description). These capabilities would allow experimentalists and theoreticians to build the models with the level of detail they need, and use them to perform *in silico* experiments. The HBP platforms would shield researchers from the complexity of the underlying technologies, provide them with the technical support they need and optimise the use of expensive supercomputing infrastructure, improving efficiency and accelerating the pace of research.

The *Neurorobotics Platform* would complement the other platforms, providing researchers with closed-loop set-ups linking brain models to simulated or physical robots. The



Figure 47: An integrated view of the brain

An integrated network of ICT platforms of the HBP

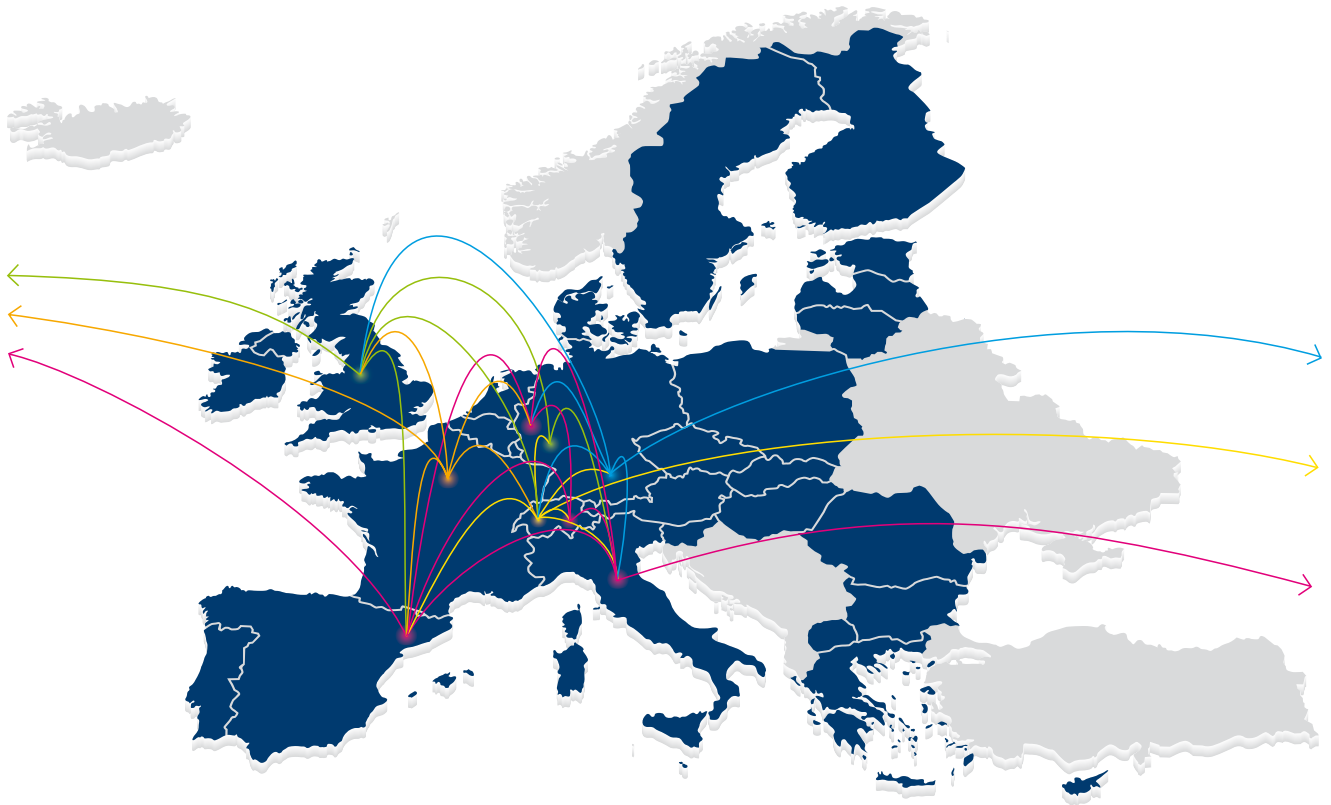


Figure 48: The HBP Platforms – physical infrastructure. Neuromorphic Computing: Heidelberg, Germany & Manchester, United Kingdom; Neurorobotics: Munich, Germany; High Performance Computing: Jülich, Germany & Lugano, Switzerland, Barcelona, Spain & Bologna, Italy; Brain Simulation: Lausanne, Switzerland; Theory Institute: Paris, France

platform would allow researchers to test theoretical models of behaviour in *in silico* experiments and to dissect the neuronal circuitry responsible.

An integrated multi-level understanding of brain structure and function

Modern neuroscience knows a lot about the individual components of the brain and has produced rigorous descriptions and conceptual models of many forms of cognition and behaviour. However, there is no agreement among neuroscientists about appropriate levels of explanation and description: for some, behaviour and cognition are best explained in terms of high-level cognitive architectures; others argue for a bottom-up approach in which they emerge from the complex dynamics of lower-level molecular and cellular systems.

The HBP – with its multi-scale modelling capabilities – would enable researchers to follow either approach – better still, to move towards a synthesis. The data, tools and platforms created by the project would enable the ambitious research programmes necessary to achieve an integrated multi-level understanding of the relationship between brain structure and function. Researchers would use the project's capabilities to explore the level of description that best explains particular electrophysiological, cognitive and behav-

oural phenomena. This would prepare the way for subsequent work to characterise the function implemented by the different levels of biological organisation, opening the door to a completely new understanding of the relationships between brain structure and function.

One of the most important issues that the HBP would make it possible to address is *learning and adaptation*: the way the brain changes its structure in response to its environment. The HBP would allow researchers to systematically trace plastic changes in model neurons and synapses as robots learn experimental tasks. Progress in this area would have a very large impact on neuroscience and would represent a major step towards brain-like information systems and robots.

A related issue is *memory*. Neuroscience has made great progress in understanding the molecular and synaptic mechanisms underlying neuronal plasticity. What it has not succeeded in achieving, to date, is an explanation of the mechanisms linking neuronal plasticity to animal and human memory. The HBP ICT platforms would help scientists to study how exposure to environmental stimuli changes the organisation and behaviour of model neuronal circuitry incorporating activity-dependent plasticity. The moment an HBP neurorobot recognises a door hiding a reward, scientists will be able to map out the way its memories are encoded and the molecules, cells, connections, and brain

regions involved. This would be a fundamental breakthrough. Better understanding of the biology of memory would make it easier to understand how it breaks down in dementia, and guide the search for new drug targets. Better theoretical understanding would help engineers to design novel information systems with memories based on the same principles as the brain. Such systems have the potential to change the face of computing.

A third vital issue is the *neural code*. Today, there is no commonly accepted theory of the way neurons code information. Theory has generated many options, but testing them is difficult. Experimental neuroscience has constructed ever more sophisticated probes, but is limited by its inability to simultaneously monitor large numbers of molecules, neurons and synapses. By contrast, the HBP's strategy of linking brain models to robots would enable researchers to study the codes used to transmit information between elements, generating predictions that could then be validated against biological experiments. Better understanding of the neural code would be a major step forward not only for neuroscience but for ICT, which is finding it increasingly hard to move data through massively parallel architectures.

Finally, HBP data and platforms would make a major contribution to high-level theories of the way the brain processes information. The HBP would provide new tools for theoreticians wishing to operationalise and test their hypotheses. Is it legitimate to view the brain as an information processing machine? If so, how can we characterise the computations it performs and how are they implemented? What aspects of brain structures are essential for their function? What are the right methodologies to bridge the gap between levels and scales, from single channels and genes, to large neuronal networks? How do emotions, thoughts, planning, language, consciousness and behaviour emerge from the underlying micro-mechanisms? *In silico* experiments would provide new opportunities to respond to these questions.

Finding answers is a precondition, not just for understanding the brain, but for any form of brain-inspired ICT.

Consciousness and awareness

One of the most difficult questions facing modern science and modern thought is consciousness [200]. What are the neuronal mechanisms and architectural principles that allow humans and other animals to transcend unconscious, basic processing as it occurs in reflexes and lower animals, and to experience the world consciously? In recent years, experimentalists have developed ingenious experimental paradigms such as masked priming [201] and binocular rivalry [202] to explore so-called Neural Correlates of Consciousness and to investigate the mechanisms responsible for the loss of consciousness in sleep, anaesthesia and neural lesions. This work has given rise to many theoretical models of consciousness including global workspaces [203], recurrent neural processing [204], and neural synchronisation [205]. To date, however, it has not been possible to design experiments providing convincing evidence in favour of a specific model.

The HBP platforms would provide an opportunity to take models that already exist, or to develop new ones, and to perform *in silico* experiments testing the models. Such experiments could potentially lead to fundamental breakthroughs.

Medicine

Clinical and pharmacological research

Just as neuroscience needs to understand how different levels of brain organisation contribute to the functioning of the healthy brain, so medical researchers need new insights into the way they break down in disease. Supplemented by the

The economic burden of brain disorders in Europe




Figure 49: Total cost of brain disease for different European countries [1]

Federating different sources of clinical data

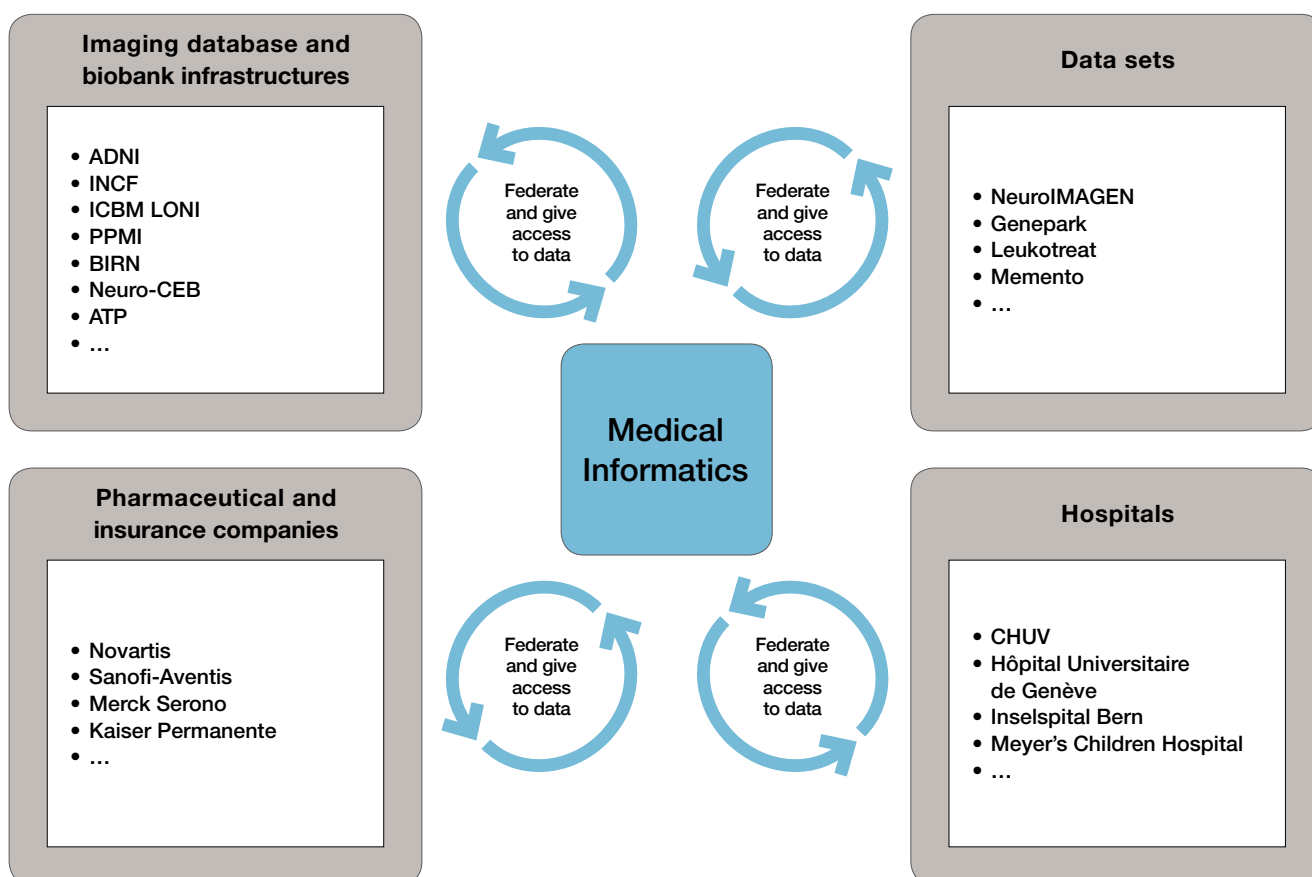


Figure 50: Medical Informatics in the HBP: federating different sources of clinical data; data mining; deriving biological signatures of disease; delivering services for personalised medicine

data and analysis tools made available by the *Medical Informatics Platform*, brain simulation could contribute to both goals, potentially transforming our understanding of brain disease, and making fundamental contributions to diagnosis and treatment.

The *Medical Informatics Platform* would federate data from hospitals and create tools allowing researchers to identify biological correlates of disease at every possible level of biological organisation, from genetics to the large-scale patterns of neuronal activity measured by imaging and EEG.

Researchers would use the HBP's data and tools to identify biological markers of disease, at different levels of biological organisation. The availability of such markers would lead to new biologically grounded classifications of disease, making it easier to identify drug targets and to select patients for participation in clinical trials. Pilot projects organised by the project would focus on autism, depression and Alzheimer's disease. Open calls would encourage use of the platform to research other diseases using a similar approach.

Many neurological and mental disorders (e.g. schizophrenia, Alzheimer's) are progressive diseases that cause severe, possibly irreversible damage, before they are diagnosed. In these cases, biological markers could make it possible to identify early stage disease, opening the prospect of early treatment to slow or halt disease progression. Even relatively minor improvements

would have a huge impact on health provider costs and on the wellbeing of patients and their families.

Simulation as a means to understand diseases of the brain

The discovery of biological signatures for a disease would inevitably suggest hypotheses of disease causation. The *Brain Simulation Platform* would allow researchers to test these hypotheses. This kind of research could provide fundamental insights into the role of different levels of biological organisation in brain disease, suggesting potential targets and strategies for treatment.

Simulation-based pharmacology

Selecting drug targets in the central nervous system involves a strong element of inference. As a result, candidate drugs for brain disease have a higher failure rate than those for other pathologies (e.g. cardiovascular disease). The tools created by the HBP would accelerate the rational design of drugs and treatment strategies, shortening and reducing the cost of the drug development cycle, lowering barriers to investment and encouraging the development of new, more effective treatments. This is a goal of fundamental importance for European health providers, European pharmaceutical companies and European society.

Impact of brain disorders in Europe




Figure 51: Prevalence and cost of brain disorders (calculated from data in [1])

Personalised medicine

Disease progression and responses to treatment vary enormously among individuals. The HBP would collect clinical data from many different individuals, from different ethnic groups, subject to different environmental and epigenetic influences. This would make it possible to compare data for individual patients. In cancer, similar methods have been used to predict the effectiveness of specific treatments in individual patients [206], contributing to the development of personalised medicine. The discovery of reliable biological signatures for psychiatric and neurological disorders would make it possible to adopt similar strategies for the treatment of these disorders, improving the effectiveness of treatment. As with other aspects of the HBP's contribution to medicine, even small improvements would have a very large impact.

Future computing technology

High performance computing

The development of supercomputing to the exascale and beyond will depend on input from user communities that will demand different classes and configurations of machine for different applications. One of the most important of these applications is simulation in the life sciences, an area in which high performance computing has yet to realise its full potential. Many of the requirements of the HBP – effective programming models, management of very large volumes of heterogeneous data, very large memory space, multi-scale simulation, interactive visualisation, interactive

steering of simulations – are common to many of the life sciences. The HBP would work closely with manufacturers to configure a supercomputing infrastructure that meets these requirements, setting the agenda for coming generations of machine.

One of the HBP's greatest needs would be to manage and analyse very large volumes of unstructured data from neuroscience experiments, simulation and the clinic. The volumes of data involved would be so large that current strategies for handling the data (generation of data on one machine, analysis and visualisation on another) and current tools (database management systems designed for business applications) would no longer be applicable. The HBP would help to meet these challenges, developing middleware that provides data base management, simulation, analytics and visualisation simultaneously on the same machine. The software framework developed by the HBP would make it possible, for the first time, for scientists to control (*steer*) data exploration, model building and simulations, interactively through a visual interface. These new techniques would have a major impact on all areas of research and engineering that use supercomputers to process large volumes of data, in particular in the life and environmental sciences.

Finally, HBP work in neuromorphic computing and insights into the brain coming from the HBP would contribute directly to the development of brain-inspired supercomputing.

Neuromorphic technologies developed by the HBP would allow the implementation of specialised analogue, digital and hybrid *boosters* for specific computational tasks. The new systems, which would provide major increases in performance and reductions in power consumption, would initially be used for brain simulation. However, they would also have applications in many other domains.

The HBP would also make broader contributions to high performance computing. In particular, studies of the neural code would contribute to the development of brain-inspired communications protocols with the potential to dramatically improve the efficiency of internal communications in high performance computing systems – currently a major source of energy consumption. Other results from the HBP would facilitate the development of novel approaches to information storage and retrieval, pattern recognition and probabilistic inference.


Taken together these developments would help transform supercomputing from an expensive, niche technology, with a handful of scientific and military applications, to a major force in 21st century ICT.

Neuromorphic computing systems


In recent years, ICT has made great progress in successfully addressing once intractable problems such as driving automatically through city streets, understanding complex queries in natural language, and machine translation. However, there are still very large areas of the economy and of daily life in which ICT has had only a limited impact. We do not use computers to clean our homes and streets, to diagnose our

A paradigm shift in future computing


Biological Brain



Numerical Simulation



Numerical Neuromorphic



Physical Neuromorphic




Figure 52: Key characteristics of the brain, conventional computing and neuromorphic computing

diseases or to look after us when we are sick. They cannot substitute a personal assistant, or a doctor or a teacher or a mechanic or a research technician. If ICT is to meet this kind of need, it will require capabilities completely lacking in current systems: massive parallelism, very low-power consumption, the ability to learn new skills without explicit programming.

Many computer scientists believe that the best way to create these capabilities is to imitate the brain. However, efforts in this direction have been hampered by poor integration between neuroscience and technology developers and by the lack of the necessary tools and technologies. One of the HBP's most important goals would be to move beyond this fragmentation, developing neuromorphic computing systems that bring together state-of-the-art hardware and design tools and simplified versions of HBP brain models. The *Neuromorphic Computing Platform* would offer developers access to these systems, providing them with the tools and support they need to explore new applications.

Key properties of neuromorphic technology, already demonstrated by the FACETS, BrainScaleS and SpiNNaker projects, include massive parallelism, very low power consumption, high resilience to failure of individual components and extremely fast computation.

The HBP would attempt to build systems that could acquire specific capabilities without explicit programming – overcoming another major barrier to the development of new ICT systems. Today, for example it is very difficult to programme systems to extract and categorise high-level information from noisy, rapidly varying sensor data. Neuromorphic computing could offer an efficient solution to this problem. Potential applications include computer vision for domestic and industrial robots, vehicles and industrial machinery, data mining for scientific research, marketing, and policing, real-time analysis of financial data (e.g. for fraud detection; rapid detection of market trends), and the monitoring of large-scale telecommunications, power distribution, and transport networks. Systems of this kind would be especially valuable in applications requiring low power consumption and high resilience to failure (e.g. wearable and portable devices, large-scale environmental monitoring, monitoring of hostile industrial environments).

Simple devices could be integrated into compact, low-power systems with the ability to control complex physical systems (e.g. vehicles, industrial machinery) with many degrees of freedom. Like the brain, such systems would have the ability to create implicit models of their environment, including their own actions and representations and those of other agents, to predict the likely consequences of their decisions, and to choose the action most likely to lead to a given goal. Although still far less flexible and powerful than the human brain, such systems would be able to perform tasks, difficult or impossible for current ICT. Examples include technical assistance to humans, real-time diagnostics of complex machinery, autonomous navigation, self-repair, and health monitoring.

Neurorobotics

One of the most important potential applications for neuromorphic computing would be in neurobots – robots whose controllers incorporate a model brain, implemented on a neuromorphic computing system. Neuromorphic controllers would benefit from many of the intrinsic advantages of neuromorphic technology, including the ability to acquire new capabilities without explicit programming, to process high-dimensional input streams, and to control robot bodies with many degrees of freedom. Physical implementations of such controllers could run up to ten thousand times faster than biological real time, allowing very rapid training. The *Neurorobotics Platform* would allow researchers and technology developers to transfer the resulting brain models to many-core devices, suitable for integration in physical robots, providing the know-how, hardware, and software they would need to explore these possibilities.

Some possible applications would be developed in pilot projects, others by groups from outside the HBP, others by industry partnerships. The compact size, high resilience and low power consumption of neuromorphic controllers would facilitate the development of a broad range of applications with a potentially very large social, industrial and economic impact. The HBP would help European industry to drive the technology instead of becoming a user and importer of technologies developed elsewhere.

The European economy and industry

Improvements in healthcare

As already reported, the cost of brain disease to the European economy has been estimated at nearly Eur 800 billion per year, accounting for 25% of the total direct cost of healthcare (costs borne by national health services, insurance companies, and patient families) and a very considerable indirect cost (lost working days for patients and their carers). Given the huge figures involved, even small improvements in treatment (e.g. improvements that delay cognitive decline in neurodegenerative disease) would produce large benefits for the European economy. If the HBP led, directly or indirectly, to effective prevention or cures for common neurological or psychiatric diseases, the economic implications would be enormous.

The pharmaceutical industry

Rational drug design requires a genuine multi-level understanding of disease mechanisms. In the case of brain disease, we do not have this understanding. This explains why pharmaceutical companies worldwide have been reducing their investment in brain research, which up to now has been less successful and less profitable than research in other fields of medicine. During the HBP-PS, major pharmaceutical companies informed us that they were strongly interested in the HBP's approach to disease and drug simulation and that they

would be willing to join future collaborative projects to test their potential in an industry setting.

Pharmaceuticals are already one of the sectors of industry in which Europe is most successful. In the HBP's new approach to rational drug design, European researchers would develop key modelling and simulation know-how and IP. The HBP thus has the potential to improve the competitiveness of the European pharmaceutical industry in what is likely to become one of the largest segments of the world pharmaceutical market.

New markets for high performance computing

The current market for high performance computing is small, and creates few economic incentives for innovation. The HBP would help to increase demand through the development of new applications that exploit the interaction and visualisation capabilities developed in the project, creating new applications for HPC in science and industry, and potentially in services for consumers. The technologies developed by the project would provide the ideal technological support for advanced services targeting industry and consumers. In this way the HBP could set in motion a virtuous circle in which massive increases in demand for high performance computing, lead to economies of scale and falling costs, and falling costs further boost demand.

Many of the new capabilities developed by the HBP would depend not on basic supercomputing hardware – which is largely developed by manufacturers from outside Europe – but on advanced software – an area in which Europe already has a strong competitive advantage. In this way, the HBP would contribute to Europe's ability to compete in what is likely to be a rapidly growing segment of the world ICT industry.

Neuromorphic computing and neurorobotics – completely new ICT

Neuromorphic computing represents a new paradigm for ICT, opening the prospect of compact, low-power systems, with a potentially unlimited range of applications, inaccessible to current technology. European research projects have already played a very important role in developing the necessary concepts, hardware and design tools. The HBP would integrate this work with research on brain modelling and simulation, giving European laboratories and companies a critical competitive edge in what is likely to become one of the most important ICT of coming decades.

Since the Industrial Revolution, the most important long-term trend in the world economy has been the mechanisation of work processes that were previously performed manually. To date the main sectors of the economy affected have been manufacturing, farming, mining, and administrative services. Neurobotic systems with neuromorphic controllers would make it possible to intensify automation in these sectors and to extend it to sectors where automation has made less progress – for instance in services and in the home. Though the first systems using neuromorphic controllers would probably

Value of major ICT markets




Figure 53: Annual world sales for industries relevant to the HBP [207-221]

be limited to relatively simple tasks, improvements in their perceptual, motor and cognitive capabilities would allow them to take on more complex work.

The initial impact is likely to involve rapid increases in productivity and parallel reductions in costs. Just as the mechanisation of manufacturing and farming enormously increased the availability of manufactured goods and food, so the deployment of robots in services and the home would vastly increase the availability of affordable services. Unlike the rise of large-scale manufacturing and farming, the new trend would be ecologically benign: a robot-controlled machine would consume no more energy than a manually controlled model. In many cases, there would also be an improvement in the quality of service.

Novel services for large populations imply the emergence of new providers and new industries to build the necessary technology. In other words, the new robotics has the potential to create completely new productive activities – contributing to economic growth. This is obviously important for European industry and for the European economy. As with previous waves of mechanisation, the new technology would change patterns of employment. While the replacement of human labour with machines (e.g. in traditional services, construction, transport etc.) would destroy employment, industries spawned by the new technology would create new jobs. The final impact would depend on political decisions and economic processes, impossible to predict in advance. What is certain is that the effects would be profound. The HBP would contribute to public debate and to academic studies of these issues through its Society and Ethics Programme.

Society and ethics

Responsible innovation

The HBP should follow a policy of Responsible Innovation, taking account of issues arising from research itself (animal experimentation, use of clinical data, ownership allocation of public spending), its potential applications (new techniques for diagnosing and treating brain disease, new classes of computing system and intelligent machine, potential military applications), and its philosophical and conceptual implications.

HBP research

Animal research

As a very large-scale project, with a significant neuroscience component, the HBP would require a large number of animal experiments – primarily with rodents. Obviously, these experiments would comply with relevant regulations and would be submitted for approval to the Institutional Review Boards; none should go beyond the normal practice of neuroscience labs around the world. Nonetheless, we recognise that this aspect of the project is a potential source of public concern.

A number of aspects of the project should mitigate this concern.

- The HBP's effort in neuroinformatics, its support for the INCF, and its collaboration with other international neuroscience initiatives would encourage a culture of data sharing, making it easy for labs to identify data which is already available and reducing duplicate animal experiments.
- Predictive neuroinformatics makes it possible to predict parameter values where experimental data is not available. Although establishing and validating predictive methods requires animal experimentation, such techniques have the long-run potential to reduce the need for experiments.
- Biologically detailed models and simulations have emergent properties that are difficult to predict from knowledge of their individual components. Once animal experiments have shown that a model successfully reproduces a certain category of measurement, researchers would be able to make the measurements in the model rather than in animals.
- The project's modelling and simulation capabilities would be especially useful for preliminary drug screening. While this kind of screening would never replace animal drug testing, it would make testing more efficient, excluding drugs that are unlikely to be effective or with predictable adverse effects.

The HBP would allow clinical researchers to gain a new understanding of brain disease, design new diagnostic tools and develop new treatments. We therefore argue that the potential contribution to human wellbeing and the potential long-term reductions in animal experimentation amply justify the project's limited and responsible use of animals.

Human studies and use of clinical data

The HBP would use of large volumes of clinical data, to inform the project's modelling effort, and to gain better understanding of brain disease. This raises the issue of how to protect patient data. Obviously the HBP would apply the best available data protection practices and technology. However, we recommend that it should supplement these techniques with new strategies allowing researchers to query hospital and other clinical data remotely, obtaining anonymous, aggregated results that do not reveal sensitive data about individual patients. These strategies are likely to be valuable not just for the HBP but also for other medical research requiring analysis of patient data on multiple sites.

A second issue with clinical data is informed consent. Some issues, such as the definition of procedures for patients who cannot express consent, are well known. However, the HBP would also pose additional problems. As an open-ended project, with far-reaching goals, the project would not be able to define in advance all possible ways in which it may wish to use patient data. This is an issue affecting many areas of large-scale clinical research and has given rise to a broad-ranging debate on open consent. Scientists and ethicists working in the HBP should participate in public and academic debate on this issue.

Ownership and public access to knowledge

Neuroscience data, clinical data and basic research tools, developed by the HBP, should become community resources freely accessible to scientists around the world. At the same time, the HBP Consortium and its members should protect and commercially exploit specific technologies, tools, and applications. This strategy would be similar to the approach adopted by other large-scale science projects. We do not expect that they would give rise to fundamental ethical or policy objections.

Allocation of public resources

In a period of financial crisis and cuts in research funding, any research proposal has an ethical obligation to demonstrate that it is better to invest in this way than to spend the money outside science, or on other forms of research. The partners in the HBP-PS believe that investment in the HBP would be fully justified by the potential benefits for European society and the European economy, described earlier in this report.

Research outcomes

Clinical applications

Some of the earliest results of HBP research would be in the clinic. Contributing to the fight against brain disease is an obvious ethical good. It nonetheless has a number of ethical, social and legal implications.

New HBP diagnostic tools would be part of a secular trend towards early diagnosis of disease. In cases where early diagnosis facilitates treatment and cure this is ethically unproblematic. However, in cases where there is no effective treatment (e.g. Huntington's disease), patients may demand their right *not to know*. It is ethically important that society finds ways of respecting this right.

Another key issue is the impact of effective early diagnosis on insurance-based health services, whose viability depends on the unpredictability of a large proportion of serious illness. Early diagnosis of highly prevalent brain diseases could seriously undermine current models of health insurance.

Finally, there can be little doubt that many of the diagnostic tools and treatments derived from HBP research would be expensive. Savings in the cost of care are likely to be far higher than the cost of early diagnosis and treatment. However, the introduction of new and expensive technologies could exacerbate inequalities in access to health care. No one would seriously argue that this is a reason to stop development of new medical technologies. It is nonetheless important that there should be an open debate on the way the new technologies are deployed, used and paid for.

Neuromorphic computing and neurorobotics

HBP work in neuromorphic computing and neurorobotics would lay the foundations for a new generation of computing systems and machines with cognitive capabilities absent in current technology, including a degree of autonomy and an ability to learn. This raises the issue of legal liability when a neuromorphic or a neurobotic system damages humans or their property – a topic already raised by the advent of autonomous vehicles. The HBP Society and Ethics Programme should contribute to this debate.

A larger, but longer-term risk is that the impact of the new technologies on employment. In some cases neuromorphic and neurobotic technologies could replace humans in tasks that are dirty, dangerous or unpleasant. It is possible, furthermore, that the new machines would provide citizens with services (e.g. help with domestic chores) that are currently unavailable or unaffordable. However, the widespread use of the new technologies would also imply a major shift in current patterns of employment and there is no guarantee that the creation of new jobs in some sectors of the economy would be sufficient to counterbalance losses in other sectors. As in the case of medical innovation, it would not be ethically justifiable to abandon technical innovation because of the risk of disruption to current economic and social structures. Again, however, these changes would need to be governed. This should be another area for debate in the HBP Society and Ethics Programme.

Military and police applications

Neuromorphic and neurobotic technologies have obvious applications in autonomous or semi-autonomous weapons systems, and as controllers for such systems. As has been pointed out in debates on military drones, the deployment of such systems raises delicate issues of morality and of international criminal law. Other questionable applications might include deployment of cognitive technologies for automated mass surveillance (analysis of images from CCTV cameras, automated transcription, translation and interpretation of phone calls, automated analysis of emails). In all these cases, public debate is essential, ideally before the technologies become available.

Philosophical and conceptual issues

Without a brain, human beings would know nothing, feel nothing, and experience nothing; they would have no subjective experience at all. Knowledge of the brain may not be enough to understand what this experience means, but it is relevant. By contributing to a better understanding of the biological mechanisms underlying decision-making, knowledge, feelings, and values, the HBP would inevitably influence conceptions of what it means to be a conscious, sentient being with a subjective view of the world. In this process, old philosophical debates would reemerge in the new light

offered by HBP advances, and new scientific issues would be philosophically analysed. What does it mean to simulate the human brain? Can the human brain be conceived independently of its context and external environment? What are the implications of the knowledge generated in HBP for our socially crucial notions of self, personhood, volitional control, and responsibility? Whatever the results of the project, they would profoundly influence current beliefs about the human mind, identity, personhood, and our capacity for control. Scientific and philosophical research in the HBP should contribute to understanding this influence and help its integration in society.



Figure 54: What does it mean to simulate the human brain?

Conclusions

In this report, we propose a FET Flagship dedicated to brain research and its applications. We call the project the *Human Brain Project*.

We suggest that the main obstacle that has prevented brain research from achieving its goals is the fragmentation of research and research data. We go on to argue that the convergence between ICT and biology has reached a point where it can allow us to achieve a genuine multi-level understanding of the brain. Such an understanding, we point out, provides the key to treating the brain's diseases and to building computing technologies with brain-like intelligence. Obviously such ambitions are too large to be achieved by any single project, including the Human Brain Project. We argue, however, that the HBP's ICT platforms would act as a catalyst, accelerating research and triggering an avalanche of additional public and industrial funding. In brief, the HBP has the potential to set off a wave of research, going far beyond the scope of the original project.

Building the HBP platforms is technically feasible. The supercomputing technology we require is developing rapidly, with each step laid out in well-established industry roadmaps. Potential partners in the HBP have already tested prototypes of key technologies for neuroinformatics, brain simulation, medical informatics, neuromorphic computing and neurorobotics. There is no doubt that Europe has the necessary know-how and resources. Realising this vision will require unprecedented collaboration on the scale of a FET Flagship.

The HBP would provide enormous value to academia, industry and society in general.

Scientifically, the HBP platforms would provide neuroscientists with the tools to integrate data from different sources, to identify and fill gaps in their knowledge, to prioritise future experiments and to trace intricate causal relationships across multiple levels of brain organisation. These possibilities would enable them to address questions inaccessible to current methods, identifying the complex cascades of events leading from genes to cognition and back, studying the biological mechanisms responsible for perception, emotions, and decision-making, and revealing the principles linking brain plasticity to learning and memory. HBP tools would even open new vistas for research into the biological mechanisms of human consciousness.

From a medical point of view, the HBP offers another revolution; a shift of paradigm from symptom and syndrome-based classifications of brain disease to a new

understanding grounded in biology. New classifications of disease would make it possible to diagnose disease at an early stage, to develop personalised treatments, adapted to the needs of individual patients and to improve chances of recovery for the afflicted. It would also reduce the economic burden on European health services, currently predicted to rise unsustainably with the rising age of the population. Given the hundreds of millions of people affected and hundreds of billions of Euros of associated costs, even small improvements would cover the cost of the HBP many times over.

These advantages for citizens are paralleled by advantages for the pharmaceutical industry. As reported earlier, pharmaceutical companies are withdrawing from research on brain disease, due to high costs and high numbers of failed clinical trials. The HBP could help to reverse this trend, speeding up drug development, cutting costs and improving success rates. Several large European pharmaceutical companies have already declared their clear intention to participate in HBP research and financially support the development of the platforms as soon as pilot projects have demonstrated their value. If this research helps them to pioneer new drugs for brain disease, they could tap a potentially enormous world market.

In neuromorphic computing, brain-inspired technologies, developed by the HBP, offer the key to managing and exploiting the noisiness and unreliability of nano- and atomic level components, overcoming fundamental limits of current computing technologies. Combined with brain-inspired techniques of data transmission, storage and learning, these technologies will make it possible to build low-cost, energy-efficient computers, ultimately with brain-like intelligence. These developments would add a completely new dimension, not just to computing, but to a broad range of 21st century technologies. A large number of companies have declared their interest in helping the HBP to realise this goal. Such systems would not replace current computing technologies but could play a complementary, equally important role, enabling completely new applications. This is a vast area of technology that Europe can lead.

The supercomputing industry will draw similar benefits. In hardware, hybridisation with neuromorphic computing would offer large improvements in performance and significant reductions in power consumption. In software, techniques of interactive supercomputing, also developed by the project, would enable completely new forms of

interactivity and visualisation, spawning radically new services, for science (“virtual instruments”), industry (simulation as a tool for engineering) and ultimately for consumer markets (advanced tools for interactive communication, education, entertainment etc.). Together, these innovations would reduce the cost and ecological impact of supercomputer power consumption, while simultaneously stimulat-

ing demand. The end result would be a major expansion of the market for European-made technologies.

In summary, the HBP is vital, timely and feasible, offering enormous potential benefits for European science, European citizens, European industry and Europe’s strategic position in the world. On these grounds, we argue that the HBP constitutes an ideal FET Flagship.



Glossary

A

Action Potential

A short-lasting electrical event in which the membrane potential of a cell rapidly rises and falls, following a consistent trajectory of depolarisation and hyperpolarisation.

ADNI

The Alzheimer's Disease Neuroimaging Initiative – a major NIH-funded initiative to compare rates of change in cognition, brain structure and biomarkers in healthy elderly people, patients with Mild Cognitive Impairment and patients with Alzheimer's disease. The project has made its imaging data freely available to the community, creating a model for HBP work in Medical Informatics.

Artificial Neural Network

A mathematical model or computational model inspired by the structure and/or functional aspects of biological neural networks.

Axon

A long projection of a neuron that conducts electrical impulses away from the principal cell body.

B

Batch mode

A way of executing one or more programmes on a computer in which the user defines the initial input data and the computer generates an output without any further interaction.

Biophysical

Refers to methods from the physical sciences, used in the study of biological systems.

Blue Brain Project

An EPFL project, launched in 2005 with the goal of creating the workflows and tools necessary to build and simulate brain models. As proof of concept, the project has successfully built and simulated a cellular-level model of the rat cortical column.

BlueGene

IBM supercomputer. The BlueGene/P

used in the EPFL Blue Brain Project is a massively parallel, tightly interconnected machine with 16384 processors, 56 Teraflops of peak performance, 16 TeraByte of distributed memory and a 1 Petabyte file system. The Blue Brain team provides enough computing power to simulate at least 60 rat cortical columns.

Booster

A special-purpose module in super-computer architecture, used to boost performance for a specific class of computations.

Brain atlas

A work of reference (e.g. the Allen Mouse Atlas), often available as an online public resource, showing how one or more data sets (e.g. gene expression data) map to specific regions and sub-regions of the brain.

Brainpedia

A community driven Wiki, proposed by the HBP-PS. The Brainpedia would provide an encyclopedic view of the latest data, models and literature for all levels of brain organisation.

BrainScaleS

An EU-funded research project that integrates *in vivo* experimentation with computational analysis to investigate how the brain processes information on multiple spatial and temporal scales and to implement these capabilities in neuromorphic technology.

C

Cable Theory

Mathematical models making it possible to calculate the flow of electric current (and accompanying voltage) assuming passive neuronal fibres such as axons and dendrites are cylindrical cable-like structures.

Connectome

The complete connectivity map between neurons, including the locations of all synapses.

Connectomics

The study of the connectome.

Cortical column

A basic functional unit of the neocortex organised as a densely interconnected column of neurons traversing all six layers.

D

Dendrite

The branched projections of a neuron that conduct electrochemical signals received from other neurons to the soma of the principal neuron.

Diffusion Tensor Imaging

A technique that enables the measurement of the restricted diffusion of water in tissue in order to produce neural tract images. It also provides useful structural information.

Division

Divisions would be the main scientific organisational units in the Human Brain Project. Each division would cover a specific discipline (e.g. cellular and molecular neuroscience, brain simulation, ethics).

DTI

See Diffusion Tensor Imaging.

E

ECoG

Intracranial electro-corticogram: a technique in which electrodes are placed directly on the exposed surface of the brain to record electrical activity.

EEG

Electroencephalography: the recording of electrical activity on the surface of the scalp. EEG measures voltage fluctuations resulting from ionic current flows within the neurons of the brain.

Electrophysiology

The study of the electrical properties of excitable biological cells and tissues.

ESS

See Executable System Specification.

Exascale

Refers to a supercomputer with a performance of 10^{18} flops. The first computers with this level of performance

are expected to become available during the second half of this decade.

Executable Systems Specification

An engineering approach to large-scale system design in which specifications are implemented as a complete software model of the device under construction. The ESS approach makes it possible to verify the hardware design without building a physical system.

F

FACETS

A European research project (2005-2010) that pioneered an integrated workflow for neuromorphic computing, leading from neurobiology and brain modelling to neuromorphic hardware.

flops

Floating Point Operations per Second. A measure of computer performance. The largest current supercomputers have a performance in the order of Petaflops (10^{15} flops). Exascale supercomputers planned for the end of the decade would have a performance in the order of exaflops (10^{18} flops).

fMRI

See Functional Magnetic Resonance Imaging.

Functional Magnetic Resonance Imaging

An MRI procedure that measures brain activity by detecting functional changes associated with changing blood flow.

G

Glia

Non-neuronal cells that maintain homeostasis, form myelin, and provide support and protection for neurons in the nervous system.

H

High performance computing

The use of parallel processing to run an applications programme efficiently, reliably and quickly. The term HPC is sometimes used as a synonym for supercomputing.

High throughput

An automated technology making it possible to generate large volumes of data at high speed and low cost, often from multiple samples.

Hodgkin and Huxley Model

A set of differential equations describing an electrical circuit model for the non-linear dynamics of ion channels and the cell membrane of neurons.

I

In silico

A process or an experiment performed

on a computer or via computer simulation.

In vitro

Studies in experimental biology conducted using components of an organism that have been isolated from their usual biological context.

In vivo

Studies using a whole, living organism as opposed to a partial or dead organism.

INCF

See International Neuroinformatics Coordinating Facility.

International Neuroinformatics

Coordinating Facility

An international science organisation whose purpose is to facilitate worldwide cooperation of activities and infrastructures in neuroinformatics-related fields.

Ion channel

Proteins controlling the passage of ions through the cell membrane. Ion channels are targets for neuromodulatory systems and for drugs. The distribution of ion channels determines the electrical behaviour of the cell.

iPSC

Induced Pluripotent Stem Cell: a type of stem cell that can be used to generate neurons and other kinds of cell for use in research.

L

Localiser

A (usually simple) task used in conjunction with fMRI to characterise the neuronal circuitry responsible for a specific cognitive or behavioural capability.

M

Machine Learning

Refers to techniques allowing a computer to discover statistical regularities in a sample of data (e.g. clinical data for patients with known diagnoses) and to exploit these regularities (e.g. by suggesting a diagnosis for a patient with an unknown pathology).

Magnetic Resonance Imaging

A medical imaging technique allowing the visualisation of detailed internal structures. Nuclear magnetic resonance (NMR) is used to image nuclei of atoms inside the body.

MCELL

A widely used simulator from the Computational Neurobiology Lab, SALK Institute, USA. Mcell is used in reaction diffusion simulations of molecular interactions.

Mechanistic

Refers to an explanation that identifies the causal chain of physical or chemical events leading from an initial cause (e.g. a gene defect) to its consequences (e.g. a change in behaviour). In clinical research, knowledge of such cascades is a precondition for rational drug design.

Microcircuit

A neural circuit lying within the dimensions of the local arborisations of neurons (typically 200–500 μm).

Molecular Dynamics

A form of computer simulation using approximations of known physics to estimate the motion of atoms and molecules.

MRI

See Magnetic Resonance Imaging.

mRNA

A molecule of RNA, transcribed from a template DNA, that acts as a blueprint for a final protein product.

Multi-level

Refers to a description of the brain that takes account of its different levels of organisation.

Multi-scale

Refers to a simulation technique that reproduces the different levels of organisation of a complex phenomenon, switching dynamically between different levels of detail according to the needs of the simulation.

N

Neural code(s)

The code or codes used by the brain to transmit information within the brain.

Neuroinformatics

The academic discipline concerned with the use of computational tools to federate, organise and analyse neuroscience data.

Neuromorphic

Refers to a method for emulating the structure and function of neurons and neuronal circuits in electronics.

Neuromorphic Computing System

A computing system comprising a neuromorphic computing device, a software environment for configuration and control and the capability to receive input and to generate output.

Neuron

An electrically excitable cell that processes and transmits information by electrical and chemical signalling.

NEURON

A well-known environment for the empirically based simulations of neurons and networks of neurons. Developed by Michael Hines, Yale University, USA.

Neurobotic System

A robotic system comprised of a controller, a body, actuators and sensors, whose controller architecture is derived from a model of the brain.

O

Optogenetics

The combination of genetic and optical methods to control specific events in targeted cells of living tissue. Optogenetics provides the temporal precision (millisecond-timescale) needed to keep pace with functioning intact biological systems.

Organelles

Specialised subunits performing a specialised function within a cell.

Patch clamp

A technique in electrophysiology that makes it possible to record from and stimulate living neural cells. The technique uses a glass micropipette to enclose an area on the surface of the membrane or patch.

P

PET

Positron Emission Tomography: an imaging technique that produces a three-dimensional image of functional processes in the body, using pairs of gamma rays emitted indirectly by a positron-emitting radionuclide (tracer).

Petascale

Refers to a supercomputer with a performance of 10^{15} flops. In November 2011, the Japanese K computer became the first machine to achieve a peak performance of more than 10 Petaflops.

Plasticity

The ability of a synapse, a neuron or a neuronal circuit to change its properties in response to stimuli or the absence of stimuli.

PLI

See Polarised Light Imaging.

Polarised Light Imaging

An imaging technique making it possible to identify the orientation of fibres in histological sections of the brain. Often used for imaging *post mortem* samples from the human brain.

Predictive Neuroinformatics

The use of computational techniques to detect statistical regularities in the relationships between two neuroscience data sets and the exploitation of these regularities to predict parameter values where experimental measurements are not available.

Proteome

The set of all the proteins expressed by a cell.

R

Receptor

A protein molecule that receives and transmits chemical information across membranes.

Resting state

The state of the brain when not engaged in any specific cognitive activity.

RNA

A chain of nucleotides, similar to DNA, important in major steps of protein formation and encoding genetic information.

S

Sequencing

A technique to determine the primary structure or sequence of a polymer.

Simulation

The imitation or replication of a complex real-world process.

sMRI

Structural Magnetic Resonance Imaging

Snap-shot model

A model representing the multi-level structure of the brain at a given stage of biological development.

Soma

The cell body or the compartment in a cell that houses the nucleus.

Spike-timing Dependent Plasticity

A process that adjusts the strength of connections between pairs of neurons based on the relative timings of inputs and output signals.

SpiNNaker

A UK funded research project whose goal is to build neuromorphic computing systems based on many-core chips with efficient bi-directional links for asynchronous spike based communication.

STDP

See Spike timing Dependent Plasticity

Steering

Refers to interactive control of a simulation using real-time (usually visual) feedback from the simulation.

STEPS

A simulator for stochastic reaction-diffusion systems in realistic morphologies, from the Theoretical Neurobiology group, University of Antwerp, Belgium.

Supercomputer

A computer whose performance is close to the highest performance attainable at a given time.

Synapse

A structure between two neurons allowing them to communicate via chemical or electrical signals.

SyNAPSE

A research project funded by the US agency DARPA with the aim of building energy-efficient, compact neuro-morphic systems based on modern component technologies.

T

Terascale

Refers to a supercomputer with a performance of 10^{12} flops.

Theory of Mind

A person or an animal is said to have a theory of mind if it can ascribe intentions to another person or animal.

Transcriptome

The set of information required to fully represent all cDNA expressed by a cell during translation of the genome.

V

Very Large Scale Integration

The integration of very large numbers of transistors on a single silicon chip. VLSI devices were initially defined as chips with more than 10,000 transistors. Current systems may contain more than 2,000,000.

VLSI

See Very Large Scale Integration.

W

Work flow

Term used in management engineering and in computer science to describe a sequence of steps leading to a well-defined outcome.

References

1. Gustavsson, A., et al., *Cost of disorders of the brain in Europe 2010*. European neuropsychopharmacology : the journal of the European College of Neuropsychopharmacology, 2011. 21(10): p. 718-79.
2. *International Technology Roadmap for Semiconductors (ITRS)*. Available from: <http://www.itrs.net/>.
3. INCF. INCF. 2010; Available from: www.incf.org.
4. Allen Institute for Brain Science *Allen Institute for Brain Science*. 2012.
5. The FACETS Project. *FACETS - Fast Analogue Computing with Emergent Transient States*. 2011 [cited 2011 June 22]; Available from: <http://facets.kip.uni-heidelberg.de/>.
6. The BrainScales Project. *BrainScales - Brain-inspired multi-scale computation in neuromorphic hybrid systems*. 2011 [cited 2011 June 22]; Available from: <http://brainscales.kip.uni-heidelberg.de/>.
7. The ECHORD Project. *ECHORD - European Clearing House for Open Robotics Development*. 2011 [cited 2011 June 24]; Available from: <http://www.echord.info/wikis/website/home>.
8. The SpiNNaker Project. *SpiNNaker - A Universal Spiking Neural Network Architecture*. 2011 [cited 2011 June 22]; Available from: <http://apt.cs.man.ac.uk/projects/SpiNNaker/>.
9. Marques, A.C. and C.P. Ponting, *Catalogues of mammalian long noncoding RNAs: modest conservation and incompleteness*. *Genome Biol*, 2009. 10(11): p. R124.
10. Lein, E.S., et al., *Genome-wide atlas of gene expression in the adult mouse brain*. *Nature*, 2007. 445(7124): p. 168-76.
11. Bayes, A., et al., *Characterization of the proteome, diseases and evolution of the human postsynaptic density*. *Nature Neuroscience*, 2011. 14(1): p. 19-21.
12. Belgard, T.G., et al., *A transcriptomic atlas of mouse neocortical layers*. *Neuron*, 2011. 71(4): p. 605-16.
13. Boyden, E.S., *A history of optogenetics: the development of tools for controlling brain circuits with light*. *F1000 Biol Rep*, 2011. 3: p. 11.
14. Bussey, T.J., et al., *New translational assays for preclinical modelling of cognition in schizophrenia: the touchscreen testing method for mice and rats*. *Neuropharmacology*, 2012. 62(3): p. 1191-203.
15. Legenstein, R., D. Pecevski, and W. Maass, *A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback*. *PLoS computational biology*, 2008. 4(10): p. e1000180.
16. Maass, W., T. Natschläger, and H. Markram, *Real-time computing without stable states: A new framework for neural computation based on perturbations*. *Neural Computation* 2002. 14(11): p. 2531-2560.
17. Fremaux, N., H. Sprekeler, and W. Gerstner, *Functional requirements for reward-modulated spike-timing-dependent plasticity*. *J Neurosci*, 2010. 30(40): p. 13326-37.
18. Wierstra, D., et al., *Recurrent policy gradients*. *Logic Journal of the IGPL*, 2010. 18(5): p. 620-634.
19. Buesing, L., et al., *Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons*. *PLoS Comput Biol*, 2011. 7(11): p. e1002211.
20. Pechura, C.M. and J.B. Martin, *Mapping the brain and its functions. Integrating enabling technologies into neuroscience research 1991*, Washington D.C.: National Academy Press.
21. Amari, S., et al., *Neuroinformatics: the integration of shared databases and tools towards integrative neuroscience*. *J Integr Neurosci*, 2002. 1(2): p. 117-28.
22. Gupta, A., et al., *Federated access to heterogeneous information resources in the Neuroscience Information Framework (NIF)*. *Neuroinformatics*, 2008. 6(3): p. 205-17.
23. Helmer, K.G., et al., *Enabling collaborative research using the Biomedical Informatics Research Network (BIRN)*. *Journal of the American Medical Informatics Association (JAMIA)*, 2011. 18(4): p. 416-22.
24. Koch, C. and R.C. Reid, *Neuroscience: Observatories of the mind*. *Nature*, 2012. 483(7390): p. 397-398.
25. Cuntz, H., et al., *One rule to grow them all: a general theory of neuronal branching and its practical application*. *PLoS computational biology*, 2010. 6(8): p. e1000877.
26. Fries, P., *Neuronal gamma-band synchronization as a fundamental process in cortical computation*. *Annual Review of Neuroscience*, 2009. 32: p. 209-224.
27. Smith, S.M., et al., *Correspondence of the brain's functional architecture during activation and rest*. *Proceedings of the National Academy of Sciences*, 2009. 106: p. 13040-13045.
28. Doeller, C.F., C. Barry, and N. Burgess, *Evidence for grid cells in a human memory network*. *Nature*, 2010. 463(7281): p. 657-61.
29. Piazza, M., et al., *Tuning curves for approximate numerosity in the human intraparietal sulcus*. *Neuron*, 2004. 44(3): p. 547-55.
30. Fox, M.D. and M.E. Raichle, *Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging*. *Nature Reviews Neuroscience*, 2007. 8: p. 700-711.
31. Berkes, P., et al., *Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment*. *Science*, 2011. 331: p. 83-87.
32. Sigman, M., et al., *Parsing a sequence of brain activations at psychological times using fMRI*. *Neuroimage*, 2007. 35(2): p. 655-68.
33. Dux, P.E., et al., *Isolation of a central bottleneck of information processing with time-resolved fMRI*. *Neuron*, 2006. 52(6): p. 1109-20.
34. Buschman, T.J. and E.K. Miller, *Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices*. *Science*, 2007. 315(5820): p. 1860-2.

35. Duncan, J., *The multiple-demand (MD) system of the primate brain: mental programmes for intelligent behaviour*. Trends Cogn Sci, 2010. 14(4): p. 172-9.
36. Hauser, M.D., N. Chomsky, and W.T. Fitch, *The faculty of language: what is it, who has it, and how did it evolve?* Science, 2002. 298: p. 1569-79.
37. Penn, D.C., K.J. Holyoak, and D.J. Povinelli, Darwin's mistake: *Explaining the discontinuity between human and nonhuman minds*. Behavioural and Brain Sciences, 2008: p. 109-178.
38. Sereno, M.I. and R.S. Huang, *A human parietal face area contains aligned head-centred visual and tactile maps*. Nat Neurosci, 2006. 9(10): p. 1337-43.
39. Kolster, H., et al., *Visual field map clusters in macaque extrastriate visual cortex*. J Neurosci, 2009. 29(21): p. 7031-9.
40. Pallier, C., A.-D. Devauchelle, and S. Dehaene, *Cortical representation of the constituent structure of sentences*. Proceedings of the National Academy of Sciences, 2011. 108: p. 2522-2527.
41. Saur, D., et al., *Combining functional and anatomical connectivity reveals brain networks for auditory language comprehension*. Neuroimage, 2010. 49(4): p. 3187-97.
42. Nieder, A. and S. Dehaene, *Representation of number in the brain*. Annual Review of Neuroscience, 2009. 32: p. 185-208.
43. Christoff, K., et al., *Specifying the self for cognitive neuroscience*. Trends in Cognitive Sciences, 2011. 15: p. 104-12.
44. Saxe, R. and L.J. Powell, *It's the thought that counts: specific brain regions for one component of theory of mind*. Psychological Science, 2006. 17: p. 692-699.
45. Miller, K.D., J.B. Keller, and M.P. Stryker, *Ocular dominance column development: analysis and simulation*. Science, 1989. 245(4918): p. 605-15.
46. Hopfield, J.J., *Neural networks and physical systems with emergent collective computational abilities*. Proc Natl Acad Sci U S A, 1982. 79(8): p. 2554-8.
47. Gerstner, W., et al., *A neuronal learning rule for sub-millisecond temporal coding*. Nature, 1996. 383(6595): p. 76-81.
48. van Vreeswijk, C. and H. Sompolinsky, *Chaos in neuronal networks with balanced excitatory and inhibitory activity*. Science, 1996. 274(5293): p. 1724-6.
49. Brunel, N., *Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons*. J Comput Neurosci, 2000. 8(3): p. 183-208.
50. Jolivet, R., T.J. Lewis, and W. Gerstner, *Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy*. Journal of Neurophysiology, 2004. 92(2): p. 959-76.
51. Badel, L., et al., *Dynamic I-V curves are reliable predictors of naturalistic pyramidal-neuron voltage traces*. Journal of Neurophysiology, 2008. 99(2): p. 656-66.
52. Markram, H., Y. Wang, and M. Tsodyks, *Differential signaling via the same axon of neocortical pyramidal neurons*. Proc Natl Acad Sci U S A, 1998. 95(9): p. 5323-8.
53. Senn, W., H. Markram, and M. Tsodyks, *An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing*. Neural Comput, 2001. 13(1): p. 35-67.
54. Clopath, C., et al., *Connectivity reflects coding: a model of voltage-based STDP with homeostasis*. Nature neuroscience, 2010. 13(3): p. 344-52.
55. El Boustani, S., et al., *Stable learning in stochastic network states*. J Neurosci, 2012. 32(1): p. 194-214.
56. Legenstein, R., et al., *A reward-modulated Hebbian learning rule can explain experimentally observed network reorganisation in a brain control task*. J Neurosci, 2010. 30(25): p. 8400-10.
57. Destexhe, A., M. Rudolph, and D. Pare, *The high-conductance state of neocortical neurons in vivo*. Nat Rev Neurosci, 2003. 4(9): p. 739-51.
58. Destexhe, A. and D. Contreras, *Neuronal computations with stochastic network states*. Science, 2006. 314(5796): p. 85-90.
59. Maass, W., T. Natschläger, and H. Markram, *Real-time computing without stable states: A new framework for neural computation based on perturbations*. Neural Computation, 2002. 14(11): p. 2531-2560.
60. Buesing, L., et al., *Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons*. PLoS computational biology, 2011. 7(11): p. e1002211.
61. Amari, S.-I., et al., *Neuroinformatics: the integration of shared databases and tools towards integrative neuroscience*. Journal of integrative neuroscience, 2002. 1(2): p. 117-128.
62. Helmer, K.G., et al., *Enabling collaborative research using the Biomedical Informatics Research Network (BIRN)*. Journal of the American Medical Informatics Association : JAMIA, 2011. 18(4): p. 416-22.
63. Song, S., et al., *Highly nonrandom features of synaptic connectivity in local cortical circuits*. PLoS biology, 2005. 3(3): p. e68.
64. Perin, R., T.K. Berger, and H. Markram, *A synaptic organizing principle for cortical neuronal groups*. Proceedings of the National Academy of Sciences of the United States of America, 2011. 108(13): p. 5419-24.
65. Hay, E., et al., *Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties*. PLoS computational biology, 2011. 7(7): p. e1002107.
66. Ascoli, G.A., et al., *Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex*. Nature reviews. Neuroscience, 2008. 9(7): p. 557-68.
67. Khazen, G., et al., *Predicting Combinatorial Gene Expression in Neocortical Neurons*. PLoS computational biology, 2011.
68. Khazen, G., et al., *Combinatorial Expression Rules of Ion Channel Genes in Juvenile Rat (Rattus norvegicus) Neocortical Neurons*. PLoS One, 2012. 7(4): p. e34786.
69. McCullough, W.S. and W. Pitts, *A logical calculus of the ideas immanent in nervous activity*. Bulletin of Mathematical Biophysics, 1943. 5 (1943): p. 115-13.
70. Hodgkin, A.L. and A.F.A. Huxley, *A quantitative description of membrane current and its application to conduction and excitation in nerve*. Journal of Physiology (London), 1952. 117: p. 500-544.
71. Rall, W., *Core conductor theory and cable properties of neurons*, in *Handbook of physiology, cellular biology of neurons*, E.R. Kandel, J.M. Brookhardt, and V.M. Mountcastle, Editors. 1977, American Physiological Society: Bethesda, MD. p. 39-97.
72. Traub, R.D., R. Miles, and R.K.S. Wong, *Large scale simulations of the hippocampus*. Engineering in Medicine and Biology Magazine, IEEE, 1988. 7(4): p. 31-38.
73. Traub, R.D. and R.K. Wong, *Synchronized burst discharge in disinhibited hippocampal slice. II. Model of cellular mechanism*. Journal of Neurophysiology, 1983. 49(2): p. 459-71.
74. De Schutter, E. and J.M. Bower, *An active membrane model of the cerebellar Purkinje cell II. Simulation of synaptic responses*. Journal of Neurophysiology, 1994. 71(1): p. 401-19.
75. De Schutter, E. and J.M. Bower, *An active membrane model of the cerebellar Purkinje cell. I. Simulation of current clamps in slice*. Journal of Neurophysiology, 1994. 71(1): p. 375-400.

76. Obermayer, K., H. Ritter, and K. Schulten, *Large-scale simulations of self-organizing neural networks on parallel computers: application to biological modelling*. Parallel Computing, 1990. 14(3): p. 381-404.
77. Izhikevich, E.M. *Simulation of large-scale brain models*. 2005; Available from: http://www.izhikevich.org/human_brain_simulation/Blue_Brain.htm#Simulation%20of%20Large-Scale%20Brain%20Models.
78. Djurfeldt, M., et al., *Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer*. IBM Journal of Research and Development, 2008. 52(1-2): p. 31-41.
79. Morrison, A., A. Aertsen, and M. Diesmann, *Spike-timing-dependent plasticity in balanced random networks*. Neural Computation, 2007. 19(6): p. 1437-67.
80. Ananthanarayanan, R., et al., *The cat is out of the bag: cortical simulations with 10^9 neurons, 10^{13} synapses*, in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis 2009*, ACM: Portland, Oregon. p. 1-12.
81. Ananthanarayanan, R. and D.S. Modha, *Anatomy of a cortical simulator, in Proceedings of the 2007 ACM/IEEE conference on Supercomputing 2007*, ACM: Reno, Nevada. p. 1-12.
82. Bednar, J.A., et al., *Modelling cortical maps with Topographica*. Neurocomputing, 2004. 58: p. 1129-1135.
83. de Kamps, M. and V. Baier, *Multiple interacting instantiations of neuronal dynamics (MIIND): a library for rapid prototyping of models in cognitive neuroscience*. 2007 IEEE International Joint Conference on Neural Networks, 2007. 1-6: p. 2834-2839.
84. Gwaltig, M.-O. and M. Diesmann *NEST (Neural Simulation Tool)*. Scholarpedia, 2007. 2, 1430.
85. Wils, S. and E. De Schutter, *STEPS: Modelling and Simulating Complex Reaction-Diffusion Systems with Python*. Frontiers in Neuroinformatics, 2009. 3: p. 15.
86. Markram, H., *The Blue Brain Project*. Nature Review of Neuroscience, 2006. 7(2): p. 153-60.
87. Druckmann, S., et al., *A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data*. Frontiers in Neuroscience, 2007. 1(1): p. 7-18.
88. Kozloski, J., et al., *Identifying, Tabulating, and Analysing Contacts between branched Neuron Morphologies*. IBM Journal of Research and Development, 2008. 52(1/2): p. 43-55.
89. Meuer, H., et al. *TOP500 Supercomputing Sites*. 2011; Available from: <http://www.top500.org/>.
90. *International Technology Roadmap for Semiconductors (ITRS)*. Available from: <http://www.itrs.net/>.
91. Kogge, P., et al., *ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems*, 2008.
92. Ashby, S., et al., *The Opportunities and Challenges of Exascale Computing*, 2010.
93. *International Exascale Software Project (IESP)*. Available from: <http://www.exascale.org/>.
94. *European Exascale Software Initiative (EESI)*. Available from: <http://www.eesi-project.eu/>.
95. *CRESTA Project - Collaborative Research into Exascale Systemware, Tools & Applications*. Available from: <http://cresta-project.eu/>.
96. *DEEP Project - Dynamical Exascale Entry Platform*. Available from: <http://www.deep-project.eu/>.
97. *Mont-Blanc Project*. Available from: <http://www.montblanc-project.eu/>.
98. Gerstein, G.L. and B. Mandelbrot, *Random Walk Models for the Spike Activity of a Single Neuron*. Biophysical Journal, 1964. 4: p. 41-68.
99. IBM. *SyNAPSE: a cognitive computing project from IBM research*. 2011; Available from: http://www.ibm.com/smarterplanet/us/en/business_analytics/article/cognitive_computing.html.
100. Ishii, S. *Create a brain on a supercomputer to unravel the functions of the brain and nervous system*. Biosupercomputing Newsletter, 2010. 2.
101. Xin, J., et al., *Modelling Spiking Neural Networks on SpiNNaker*. Computing in Science & Engineering, 2010. 12(5): p. 91-97.
102. Markram, H., *The Blue Brain Project*. Nature Review of Neuroscience, 2006. 7(2): p. 153-60.
103. Hughes, A.J., et al., *What Features Improve the Accuracy of Clinical-Diagnosis in Parkinsons-Disease - a Clinicopathological Study*. Neurology, 1992. 42(6): p. 1142-1146.
104. Beach, T.G., et al., *Accuracy of the Clinical Diagnosis of Alzheimer Disease at National Institute on Aging Alzheimer Disease Centres, 2005-2010*. Journal of Neuropathology & Experimental Neurology, 2012. 71(4): p. 266-273 10.1097/NEN.0b013e31824b211b.
105. Visscher, K.M. and D.H. Weissman, *Would the field of cognitive neuroscience be advanced by sharing functional MRI data?* BMC Medicine, 2011. 9.
106. Stonnington, C.M., et al., *Interpreting scan data acquired from multiple scanners: a study with Alzheimer's disease*. Neuroimage, 2008. 39(3): p. 1180-5.
107. ADNI. Available from: <http://adni-info.org>.
108. Mead, C.A., *Analogue VLSI and Neural Systems* 1989: Addison-Wesley.
109. Indiveri, G., et al., *Neuromorphic silicon neuron circuits*. Frontiers in Neuroscience, 2011. 5: p. 73.
110. Millner, S., Grübl, A., Meier, K., Schemmel, J., Schwartz, M.O. *A VLSI Implementation of the Adaptive Exponential Integrate-and-Fire Neuron Model*. in *Advances in Neural Information Processing Systems*. 2010. MIT Press.
111. Schemmel, J., et al. *Modelling Synaptic Plasticity within Networks of Highly Accelerated I&F Neurons*. in ISCAS. 2007. IEEE Press.
112. Philipp, S., et al. *Interconnecting VLSI Spiking Neural Networks Using Isochronous Connections*. in IWANN. 2007. Springer LN in Computer Science.
113. Davison, A.P., et al., *PyNN: a common interface for neuronal network simulators*. Frontiers in Neuroinformatics, 2009. 2.
114. Sharp, T., T. Patterson, and S. Furber, *Distributed Configuration of Massively-Parallel Simulation on SpiNNaker Neuromorphic Hardware*, in *International Joint Conference on Neural Networks (IJCNN 2011)* 2011: San Jose, California.
115. University of Manchester. *SpiNNaker Home Page*. 2012 [cited 2012 April 18, 2012].
116. Furber, S. and A. Brown, *A Biologically-Inspired Massively-Parallel Architectures - Computing Beyond a Million Processors*, in *9th Conference on Application of Concurrency to System Design* 2009.
117. Khan, M.M., et al., *Mapping neural networks onto a massively-parallel chip multiprocessor in IEEE World Congress on Computational Intelligence* 2008.
118. Lund, H.H., B. Webb, and J. Hallam, *Physical and Temporal Scaling Considerations in a Robot model of Cricket Calling Song Preference*. Artificial Life, 1998. 4(1): p. 95-107.
119. Burgess, N., *Spatial cognition and the brain*. Ann N Y Acad Sci, 2008. 1124: p. 77-97.
120. Luque, N.R., et al., *Cerebellar-like corrective-model abstraction engine for robot movement control*. Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on, 2011.

121. Caligiore, D., et al., *TRoPICALS: a computational embodied neuroscience model of compatibility effects*. *Psychol Rev*, 2010. 117(4): p. 1188-228.
122. Burgess, N., J.G. Donnett, and J. O'Keefe, *Using a Mobile Robot to Test a Model of the Rat Hippocampus*. *Connection Science*, 1998. 10(3-4): p. 291-300.
123. Calabretta, R., *Genetic interference reduces the evolvability of modular and non-modular visual neural networks*. *Philos Trans R Soc Lond B Biol Sci*, 2007. 362(1479): p. 403-10.
124. Calne, D.B., et al., *Positron emission tomography after MPTP: observations relating to the cause of Parkinson's disease*. *Nature*, 1985. 317(6034): p. 246-8.
125. Kertesz, A., et al., *The evolution and pathology of frontotemporal dementia*. *Brain*, 2005. 128(Pt 9): p. 1996-2005.
126. Tate, S.K., et al., *Genetic predictors of the maximum doses patients receive during clinical use of the anti-epileptic drugs carbamazepine and phenytoin*. *Proceedings of the National Academy of Sciences of the United States of America*, 2005. 102(15): p. 5507-12.
127. Gutmann, A., *The Ethics of Synthetic Biology: Guiding Principles for Emerging Technologies*. *Hastings Centre Report*, 2011. 41(4): p. 17-22.
128. Beauchamp, T.L. and J.F. Childress, *Principles of Biomedical Ethics*. 6th ed2008, New York: Oxford University Press USA.
129. Guston D, P.J., Tosi J., *Anticipating the ethical and political challenges of human nanotechnologies*, in *Nanoethics: The Ethical and Social Implications of Nanotechnology* F. Allhoff, Lin, P. , Moor, J. , Weckert J. , Editor 2007. p. 185-97.
130. Stirling, A., "Opening Up" and "Closing Down.". *Science, Technology & Human Values*, 2008. 33(2): p. 262.
131. Joly, P.B., A. Rip, and M. Callon, *Reinventing innovation*, in *Governance of Innovation*, M. Arentsen, W. van Rossum, and B. Steenge, Editors. 2010, Edward Elgar.
132. Knorr-Cetina, K., *Epistemic cultures : how the sciences make knowledge* 1999, Cambridge, Ma: Harvard University Press.
133. Knorr-Cetina, K., *Culture in global knowledge societies: knowledge cultures and epistemic cultures*. *Interdisciplinary Science Reviews*, 2007. 32(4): p. 361-75.
134. Becker, H.A. and F. Vanclay, *The international handbook of social impact assessment: Conceptual and methodological advances* 2003, Cheltenham: Edward Elgar Publishing.
135. Abi-Rached JM, R., N *The birth of the neuromolecular gaze*. *History of the Human Sciences*, 2010. 23(1): p. 11-26.
136. Changeux, J.P., *Neuronal man: the biology of mind* 1997, Princeton, N.J: Princeton University Press.
137. Churchland, P.S., *The impact of neuroscience on philosophy*. *Neuron*, 2008. 60: p. 409-11.
138. Changeux J-P, R.P., *What makes us think? : a neuroscientist and a philosopher argue about ethics, human nature, and the brain* 2002, Princeton, N.J, Oxford: Princeton University Press.
139. Bennett MR, H.P., *Philosophical foundations of neuroscience* 2003, Oxford: Blackwell Pub.
140. Andreason, N., *Brave New Brain: Conquering mental illness in the era of the genome* 2001, New York: Oxford University Press.
141. Evers, K., *Personalized medicine in psychiatry: ethical challenges and opportunities*. *Dialogues in Clinical Neuroscience*. *Personalized Medicine: Prediction, Prevention, Participation*, 2009. 11(4): p. 427-433.
142. Dudai, Y., *The neurosciences: The danger that we will think that we have understood it all.*, in *The new brain sciences*, D. Reese and S. Rose, Editors. 2004, Cambridge University Press. p. 167-180.
143. Hyman, S., *A glimmer of light for neuropsychiatric disorders*. *Nature*, 2008. 455: p. 890-3.
144. Logothetis, N., *What we can do and what we cannot do with fMRI*. *Nature*, 2008. 453: p. 869-78.
145. Metzinger, T., *The Ego Tunnel: The Science of the Mind and the Myth of the Self* 2009, New York: Basic Books.
146. Vidal, F., *Brainhood, anthropological figure of modernity*. *History of the Human Sciences*, 2009. 22: p. 5-36.
147. Ramachandran, V.S., *The tell-tale brain: unlocking the mystery of human nature* 2011, London: Heinemann.
148. Haggard, P., *Human volition: towards a neuroscience of will*. *Nature Reviews Neuroscience*, 2008. 9(12): p. 934-46.
149. Spence, S., *The actor's brain : exploring the cognitive neuroscience of free will* 2009, New York: Oxford University Press.
150. Evers, K., *Neuroéthique. Quand la matière s'éveille* 2009, Paris: Éditions Odile Jacob.
151. Dehaene, S. and J.-P. Changeux, *Experimental and theoretical approaches to conscious processing*. *Neuron*, 2011. 70: p. 200-27.
152. Glannon, W., *Neuroethics*. *Bioethics*, 2006. 20(1): p. 37-52.
153. Illes, J., *Neuroethics : defining the issues in theory, practice, and policy* 2006, Oxford: Oxford University Press.
154. Levy, N., *Neuroethics: Challenges for the 21st Century* 2007 New York: Cambridge University Press.
155. Evers, K., *Toward a philosophy for neuroethics*, *EMBO Report*, special issue on neuroscience and society, 2007. p. 48-51.
156. Illes, J. and E. Racine, *Imaging or imagining? A neuroethics challenge informed by genetics*. *Am J Bioeth*, 2005. 5(2): p. 5-18.
157. Brownsword, R., *Regulating brain imaging: challenge and opportunity*. *Frontiers in Human Neuroscience*, 2009. 3: p. 50.
158. Farah MJ, I.J., Cook-Deegan R et al., *Neurocognitive enhancement: What can we do and what should we do?* *Nature Reviews Neuroscience*, 2004. 5 (5): p. 421-5.
159. Turner, D., *Neuroethics of Cognitive Enhancement*. *BioSocieties*, 2006. 1: p. 113-23.
160. Greely, H., B. Sahakian, and J. Harris, *Towards responsible use of cognitive-enhancing drugs by the healthy*. *Nature*, 2008. 456(7223): p. 702-5.
161. Abi-Rached JM, D.Y., *The Implications of memory research and 'memory erasers'*. *BioSocieties*, 2009. 4: p. 79-80.
162. Garland, B., *Neuroscience and the Law: brain, mind and the scales of justice*. 2004, New York: Dana Press.
163. Morse, S., *Determinism and the death of folk psychology: two challenges to responsibility from neuroscience*. *Minnesota Journal of Law, Science and Technology*, 2008. 9: p. 1-36.
164. Tovino, S., *The impact of neuroscience on health law*. *Neuroethics*, 2008. 1(2): p. 101-17.
165. Moreno, J.D., *Mind wars : brain research and national defense* 2006, New York: Dana Press.
166. Andersen, I.-E. and B. Jaeger, *Scenario workshops and consensus conferences: towards more democratic decision-making*. *Science and Public Policy*, 1999. 26(5): p. 331-340.
167. Joss, S. and S. Belucci, *Participatory Technology Assessment: European Perspectives*. *London: Centre for the Study of Democracy* 2002, University of Westminster.
168. Rowe, G. and L.J. Frewer, *A Typology of Public Engagement Mechanisms*. *Science, Technology & Human Values*, 2005. 30(2): p. 251-290.
169. Fiorino, D., *Environmental risk and democratic process: a critical review*. *Columbia Journal of Environmental Law*, 1989. 14: p. 501-47.

170. Willsdon, J. and R. Wills, *See through Science: why public engagement needs to move upstream* 2004, London: Demos.
171. European Commission, *Understanding Public Debate on Nanotechnologies - Options for Framing Public Policy* 2010, Directorate-General for Research Science, Economy and Society.: Brussels.
172. Callon, M., P. Lascoumes, and Y. Barthe, *Acting in an uncertain world: an essay on technical democracy* 2009: The MIT Press.
173. Pellizzoni, L., *Responsibility and environmental governance*. Environmental Politics, 2004. 13(3): p. 541-565.
174. Kjolberg, K.L. and R. Strand, *Conversations About Responsible Nanoresearch*, in NanoEthics 2011. p. 1-15.
175. European Commission, *From Science and Society to Science in Society: Towards a Framework for "Co-Operative Research"* 2006, European Commission: Brussels.
176. The European Parliament and the Council of the European Union, *Directive 2010/63/EU of the European Parliament and of the Council of 22 September 2010 on the protection of animals used for scientific purposes.*, 2010, Official Journal of the European Union. p. L 276/33.
177. Gottweis, H., *Stem Cell Policies in the United States and in Germany*. Policy Studies Journal, 2002. 30(4): p. 444-469.
178. Juengst, E.T., *Self-Critical Federal Science? The Ethics Experiment within the U.S. Human Genome Project*. Social Philosophy and Policy, 1996. 13(02): p. 63-95.
179. GeneNetwork. Available from: <http://www.genenetwork.org/webqtl/main.py>.
180. The Mouse Brain Library (MBL). Available from: <http://www.mbl.org/>.
181. NIH Human Connectome Project. Available from: www.humanconnectome.org.
182. CEA-Neurospin. Available from: <http://www.meteoreservice.com/neurospin/>.
183. The Blue Brain Project. Available from: <http://bluebrain.epfl.ch/>.
184. Forschungszentrum Jülich (FZJ). Available from: http://www.fz-juelich.de/portal/DE/Home/home_node.html.
185. Partners for Advanced Computing in Europe (PRACE). Available from: <http://www.prace-project.eu/>.
186. Swiss National Computing Centre (CSCS). Available from: <http://www.cscs.ch/>.
187. The Barcelona Supercomputing Centre (BSC). Available from: <http://www.bsc.es/>.
188. Universidad Politecnico de Madrid (UPM). Available from: <http://www.upm.es/institucional>.
189. European Commission, *She Figures 2009 - Statistics and indicators on gender equality in science*, 2009, European Commission - European Research Area: Brussels, Belgium.
190. Ruest-Archambault, E., *Benchmarking policy measures for gender equality in Science*, 2008, European Commission,.
191. *Koalitionsvertrag zwischen CDU, CSU und SPD: 17. Legislaturperiode*. 2009.
192. Ministère de l'enseignement supérieur et de la recherche *Stratégie nationale de recherche et d'innovation (S.N.R.I)*. 2012.
193. Ministero dell'Istruzione dell'Università e della Ricerca, *Programma Nazionale della ricerca 2011-2013*, 2011, Ministero dell'Istruzione dell'Università e della Ricerca,.
194. Comision interministerial de ciencia y tecnologia, *Plan Nacional de Investigacion Cientifica, Desarrollo e Innovacion Tecnologica*, 2008, Comision interministerial de ciencia y tecnologia,.
195. Medical Research Council *High level long term priorities*. 2012.
196. The European Commission, *Europe 2020. A strategy for smart, sustainable and inclusive growth*, 2010: Bruxelles.
197. Ben-Israel, I., *Personal Communication*, 2010: Israel National Council for Research and Development.
198. Schwab, M.E. and A.D. Buchli, *Drug research: Plug the real brain drain*. Nature, 2012. 483(7389): p. 267-268.
199. One Mind for Research *Patrick Kennedy: Joe Biden to headline research event*. 2012.
200. Koch, C., *The quest for consciousness* 2004: Roberts and Company Publishers.
201. Vorberg, D., et al., *Different time courses for visual perception and action priming*. Proceedings of the National Academy of Sciences of the United States of America, 2003. 100(10): p. 6275-80.
202. Leopold, D.A. and N.K. Logothetis, *Multistable phenomena: changing views in perception*. Trends Cogn Sci, 1999. 3(7): p. 254-264.
203. Dehaene, S. and L. Naccache, *Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework*. Cognition, 2001. 79(1-2): p. 1-37.
204. Lamme, V.A., *Towards a true neural stance on consciousness*. Trends Cogn Sci, 2006. 10(11): p. 494-501.
205. Engel, A.K. and W. Singer, *Temporal binding and the neural correlates of sensory awareness*. Trends Cogn Sci, 2001. 5(1): p. 16-25.
206. Borden, E.C. and D. Raghavan, *Personalizing medicine for cancer: the next decade*. Nat Rev Drug Discov, 2010. 9(5): p. 343-344.
207. Marketline, *PCs: Global Industry Guide*, 2011.
208. Marketsandmarkets.com, *Global Home health care Market worth US\$159.6 Billion By 2014*, 2010.
209. Marketline, *Networking Equipment: Global Industry Guide*, 2011.
210. Marketline, *Mobiles Phones: Global Industry Guide*, 2010.
211. Urch Publishing, *Pharmaceutical Market Trends, 2010 - 2014: Key market forecasts & growth opportunities*, 2009.
212. Reuters, *Factbox: A look at the \$65 billion video games industry*, 2011.
213. Marketsandmarkets.com, *Cloud Computing Market: Global Forecast (2010 - 2015)*. 2010.
214. Intersect360 research, *HPC Total Market Forecast for 2011 through 2015: Economic Sectors and Vertical Markets*. 2011.
215. Marketsandmarkets.com, *Global Automotive Semiconductors Industry Forecast (2010-2014)*, 2010.
216. IFR Statistical department, *World Robotics 2011 Industrial Robots*, 2010.
217. Neuroinsights, *The neurotechnology industry 2010 report*, 2010.
218. Gartner, *Market Share: All Software Markets, Worldwide*, 2011.
219. Marketsandmarkets.com, *Global Service Robotics Market (2010-12015)*, 2010.
220. RNCOS, *Bioinformatics Market Outlook to 2015*, 2010.
221. ABi Research, *Personal Robotics* 2010.

Acknowledgements

The authors gratefully acknowledge the hundreds of scientists who contributed ideas, text and criticism to this report. We express our special thanks to the members of the Central Project Office team (Christian Fauteux, Melissa Cochrane, Amanda Pingree, Marija Rakonjac, Keerthan Muthurasa, Nenad Buncic) and to the project managers in the partner organisations (Björn Kindler, Mirjam Delesky, Boris Orth, Claire Devillers, Pilar Flores-Romero, Tricia Foat, Konstantinos Dalamagkidis, Florian Röhrbein, Racheli Kreisberg, Benjamin Simmenauer, Christiane Riedl, Simone Mergui, Giovanna Santoro, Özlem Cangar).

We thank Stefan Eilemann, Eilif Muller, Martin Telefont, Fabien Delalondre, and Gabriel Mateescu from the Blue Brain Project for their valuable input to the roadmaps and the report.

We thank Geerd-Rüdiger Hoffmann, Thomas Heinis, Bernd Hentschel and Juan Hernando for their input to high performance computing.

We thank International Data Corporation (IDC) for valuable feedback on market values.

We thank Martin Curley (Intel) for valuable feedback on future supercomputing roadmaps.

We thank Haim Sompolinski, Andreas Herz and other members of the Bernstein Center for Computational Neuroscience, and Kris Verstreken and other members of the neurotechnologies task force for their valuable contribution to the initial evolution of the HBP report.

We also thank all those scientists who submitted proposals in the initial phase and who helped shape the discussions and ideas of the project.

We thank all the local administrative staff at the various institutions having helped in supportive activities and organizing meetings.

quality imaging health Society & Ethics Neuroinformatics
Multidisciplinary transformative neuroscience Multilevel
enable build insights strategy simplify dedicated vision
quality visualisation platforms aggregation resources fundamental
neurons signatures biology quality environment drugs
model spatial optogenetics breakthroughs software performance model
responsibility education standards design category fragmentation
Molecular & Cell language human memory compute
dynamics parallel technology partners benefits
Neuroscience data time network levels molecule
proteome fellowships transformation consciousness Theoretical
transcriptome framework Europe innovation algorithm
compute horizon receptors future hospitals power Neuroscience
digital Brain diseases ICT valuable sustainability
data training synapse workflows organise Theory Medical
quality Simulation organise learning Informatics
sharing plasticity closed loop analysis patterns self-adaptive
vision Neuromorphic predict analysis
clinical trails Computing knowledge international task open
process side expression genes HBP goals pharmacology
effects connect motor focus in silico simplify Cognitive
accelerate scale understanding core neuroscience
robots inspired life implement foundation standardised
Impact study emulation challenges detailed leading
functional diversity navigation Future real-time beyond
data objective demonstrate Bayesian systematic low energy
paradigms unifying Computing compare mining
debate mechanisms participants extreme regions tools quality
function electronics interactive collaborate Informatics
data management behaviour model
drugs multiscale facilitate understand virtual world data
paradigm physical devices connectome testing stochasticity
shift essential potential applications science revolutionary
capabilities hypothesis predictive complexity compute
roadmaps machine development advanced screening concerted
quality engineering integrate industry hardware
model Supercomputing service discover develop High
interdisciplinary species animals genome circuits Performance
data neurotransmitter Neurorobotics architectures Computing

