Abigail Morrison

IAS-6/INM-6
Jülich Research Centre, Jülich
Computer Science 3 - Software Engineering
RWTH Aachen
+49 (0)2461 619805
morrison@fz-juelich.de
www.fz-juelich.de/inm/inm-6/

OVERVIEW

Professor of computational neuroscience with a primary focus on functional spiking neural network models for learning and representation, and a secondary focus on the software ecosystem for high-performance neural network simulation.

APPOINTMENTS

Professor of Computational Neuroscience (W2) 2020-
Computer Science 3: Software Engineering - RWTH Aachen

- Teaching courses on neuroinspired computing

Leader of Computation in Neural Circuits Group 2012-
Insitute of Advanced Simulation (IAS-6) / Institute of Neuroscience and Medicine (INM-6)
Jülich Research Centre

- Research on fundamental computational properties, particularly learning and spike based computation paradigms, in healthy and diseased neuronal network models
- Research on high-performance simulation technology and interfaces of simulation technology to non-neuronal applications

Scientific Lead of Simulation Laboratory Neuroscience 2013-2023
Jülich Supercomputing Centre – Jülich Research Centre

- Research and support in high-performance applications for simulation and data analytics, domain specific languages for model definition, interactive visualisation and steering, neuromorphic computing

Professor of Computational Neuroscience (W2) 2012-2020
Institute of Cognitive Neurosciences – Ruhr-University of Bochum

- Teaching courses on scientific computing, simulation of biological neuronal networks, computational modelling in neuroscience

Junior Professor in Computational Neuroscience 2009-2012
Faculty of Biology and Bernstein Center Freiburg – Albert-Ludwigs University of Freiburg

- Research on synaptic plasticity and learning, hybrid event-driven/time-driven simulation algorithms

- Teaching courses on scientific computing, simulation of biological neuronal networks and cognitive neuroscience

Research Scientist 2006-2009
Computational Neurophysics Lab – RIKEN Brain Science Institute

- Research on spike-timing dependent plasticity, neuronal models of reinforcement learning, high-performance simulation technology

Postdoc 2006
Bernstein Center Freiburg – Albert-Ludwigs University of Freiburg

- Research on spike-timing dependent plasticity, efficient and accurate algorithms for solving neuronal dynamics, distributed simulation technology

EDUCATION

Postgraduate Diploma in Management of Software Projects 2008-2014
Open University, UK
Distinction; part-time study

PhD in Computational Neuroscience 2001-2006
Albert-Ludwigs University Freiburg; Faculty of Biology
Summa cum laude. "Spike-timing dependent plasticity in recurrent networks"
Academic advisors: Prof. Dr. Markus Diesmann and Prof. Dr. Ad Aertsen
Research carried out at the Bernstein Center Freiburg (2004-2006) and the Max-Planck Institute for Dynamics and Self-Organization, Göttingen (2001-2003)

MSc in Artificial Intelligence 2000-2001
University of Edinburgh, Uk; School of Informatics
Distinction. "How the basal ganglia learn to select actions"
Howe Master's Prize for best MSc in Artificial Intelligence (2001)
Academic advisor: Prof. Dr. David Willshaw

Studies of Physics with Computer Science minor 1997-2000
Georg-August University Göttingen; Faculty of Physics
Vordiplom "Sehr gut" (distinction) attained Dec. 2000

Studies of Physics and Philosophy 1995-1997
Kings College London; School of Physics

THIRD PARTY FUNDING

- JARA ERS, "Towards an integrated data science of complex natural systems", €400,000 (co-PI, 10/20 - 11/21).
- JARA ERS, "Recurrence and stochasticity for neuro-inspired computing", €106,600 (co-PI, 6/19 - 5/20).

- Helmholtz Association, “SO-092: Advanced computing architectures”, €8,851,000 (co-PI, 11/18 - 9/22).
- BMBF (Germany) / NIH (USA), “CompConnectomics”, €279,000 (co-PI, 12/15 - 11/18).
- JARA Seed Fund, “NESTML”, €250,000 (Joint Lead PI, 01/15 - 12/16).
- Helmholtz Association, “Funding of positions for top female researchers (W2/W3 programm)”, €750,000 (sole PI, 01/14 - 12/19).
- BMBF (Germany) / JST (Japan), “Neural circuit mechanisms of reinforcement learning”, €203,000 (Joint Lead PI, 04/13 - 10/16).
- EU 7th Framework Programme, “Human Brain Project”, (co-WP-Leader / Task Leader, 04/13 - 09/23)
- DFG, “KFO 219 - Subproject 9”, €264,000 (co-PI, 05/12 - 04/15).
- Baden-Württemberg Junior Professor Program, “Motor compositionality and task switching in spiking neuronal networks”, €141,000 (PI, 09/10 - 08/13).
- Neurex, “Coming back of researchers grant”, €20,000 (PI, 07/10 - 06/11).
- Jülich Research Center, ”Brain-Scale Simulations”, 12 Blue Gene rack-months, 1 year (extended yearly since 09/08)
- Helmholtz Initiative on Systems Biology, “Human Brain Model”, €4,400,000 (co-PI, 01/08 - 12/12).

CURRENT AND PAST TEACHING

- **Current courses**
  - **Neuroinspired Computing** - combined lecture/seminar series on the undergraduate/postgraduate level
- **Previous courses**
  - **Scientific Programming** - undergraduate/postgraduate practical course with project work and accompanying lecture series
  - **Simulation of Biological Neuronal Networks** - undergraduate/postgraduate practical course with accompanying lecture series
  - **Computational Modelling in Cognitive Neuroscience** - introductory seminar
  - **Computational Neuroscience I/II** - undergraduate/postgraduate lecture series with accompanying exercises (individual modules)
  - **Cognitive Psychology** - MSc level lecture series within a Cognitive Neuroscience specialisation unit
- **Curriculum development** Member of conception and design team for the Neuroscience modules of the revised MSc Biology course in Freiburg, in particular components on scientific programming and simulation (2011-2012)

- **Student supervision** Since my first appointment as professor in November 2009 I have supervised to completion five MSc theses and eleven PhD dissertations. I am currently supervising three MSc theses and five PhD dissertations.

OTHER SCIENTIFIC ACTIVITIES

- Member of steering committee and key developer for NEST (www.nest-simulator.org), a collaborative open-source project developing technology for the high-performance simulation of neural systems
- Deputy Workpackage Leader (WP 7.3) of the EU Flagship Project "The Human Brain Project" (2018-2020); Task Co-leader (T5.3) (2020 - present)
- Topic Spokesperson for "Theory, modelling and simulation" in the Helmholtz research programme "Decoding the Human Brain" 2013 - 2016; deputy spokesperson 2017 - 2019
- Workpackage leader in the Helmholtz Portfolio Theme "Supercomputing and Modeling for the Human Brain" (2013 - 2017)
- Member of JARA-HPC
- Domain co-chair for Life Sciences, PASC18
- Co-organizer of a series of workshops on HPC in Neuroscience at the Jüich Supercomputer centre, the CodeJam meetings 2009-2016 and a CNS*11 workshop: "Supercomputational Neuroscience - tools and applications"
- Associate editor for PLOS Computational Biology and review editor for Frontiers in Neuroinformatics
- Grant proposal reviewer for multiple national and international funding programmes
- Chief Scientific Officer of evolu8, a start-up developing GPU-accelerated automated classification of medical images, until the IPR were purchased by Zebra Medical Vision in 2016.

NON-SCIENTIFIC SKILLS AND ACTIVITIES

I speak English natively, German fluently, and French badly. When I am not doing science, I enjoy playing the trumpet, running medium to long-ish distances, cooking, and reading, especially science fiction.

RESEARCH PROFILE AND PLAN

In recent times machine learning has made immense progress, due partially to conceptual advances, but also to increased computer power and the availability of large bodies of labelled training data. However, there remain many tasks in which brains show vastly superior performance, for example those involving creativity, generalisation and adapting learned behaviour to new contexts. In order to address such tasks, the fundamental computational principles of the brain must be identified and formulated in abstract terms. My work towards this goal can be divided into two interconnected areas: computational neuroscience and simulation technology. In the former I aim to uncover how the functions of the brain can be realised by its hardware, specifically the constraints on structure, dynamics and plasticity. In the latter I develop high performance data structures and algorithms to enable efficient simulation of spiking neuronal networks, and related tools in the simulation ecosystem. These two areas converge in the field of neuromorphic computing.

Computational Neuroscience

Computational circuits in the brain consist of very many relatively simple units (neurons) communicating sporadically (< 10 Hz) over low information channels (synapses). Although each neuron receives thousands of inputs, the 10^9 synapses located in a local area of the cortex (1 mm^3) are distributed over 10^3 neurons, resulting in a connection probability of around 10%. This makes the brain a fundamentally different system from traditonal von Neumann devices. Many essential questions on how the brain stores and processes information remain to be answered. Progress in this area is not only important for understanding how the brain works, but also will enable us to exploit the low-energy, flexible computing properties of neural systems for future applications in science and industry.

In my group we focus on three particular aspects of this. Firstly, reinforcement learning is a model that accounts for how the brain can select actions that maximise reward and minimise punishment in novel environments. Whereas previous work (including ours) has demonstrated that spiking neuronal networks can indeed implement such algorithms (Potjans et al., 2009, 2011; Jitsev et al., 2012; Morita et al., 2016), this has only been shown for problems in which the input state space was pre-defined and cleanly separable. Our current work focuses on investigating how meaningful neuronal activity states can self-organise in the more realistic scenario that no such externally imposed structure exists. Following this, the research plan is to develop closed-loop experiments in which such self-organising neuronal networks control virtual robotic agents (see next section for corresponding technology). In this way, we can investigate the evolution of the self-organisation on the basis of (rewarded) interaction with an environment, and uncover hypotheses for the biological mechanisms responsible for generating a reward prediction error signal on the basis of such states.

Secondly, we investigate reservoir computing, a powerful conceptual framework based on a similar principle to support vector machines, that explains how the spiking activity of neurons can classify and perform computations on continuous input streams (Toledo-Suárez et al., 2014; Duarte & Morrison, 2014; Duarte et al., 2017). Our work in this area aims to uncover the relationship between structure, dynamics and computational capacity, and to explore the potential for interaction with other neural computing paradigms such as reinforcement learning, as mentioned above. For example, we are currently investigating hierarchical systems of reservoir computing networks and the role of topographic mappings, common in sensory systems, in accurate information transition. Our recent work has uncovered a variety of issues with reservoir computing that we intend to address in the coming few years. One aspect of this is metrics for computational capacity. One can always test a network empirically by measuring its performance on a particular task. However, this is computationally expensive, and not necessarily informative with respect to the performance of the network on other tasks. Some previous studies have proposed metrics to capture the *generic* computational capacity, however these have not been adequately validated. Thus, we will evaluate existing metrics on a battery of tasks of varying complexity (e.g. degrees of freedom, time dependency) and develop appropriate metrics if existing ones are insufficient. A further issue is the relationship of biological features of reservoir computing networks (e.g. recurrency, plasticity, heterogenity) to their computational capacity, which has so far only been addressed cursorily by the field. Again, this requires a systematic approach using a range of example tasks. Of particular interest are those biological features that may represent constraints for neuromorphic computing systems, see final section.

Thirdly, examining how the dynamic and computational properties of networks are affected by neurodegenerative diseases such as Alzheimer's and Parkinson's (Bahuguna et al., 2017; Bachmann et al., 2018) or under the influence of optogenetic stimulation (**?**), gives us the opportunity to gain insight in the necessary properties of the healthy brain. As a side effect, our approach has generated some novel approaches for automated diagnosis based on analysis of functional magnetic resonance imaging data. Our current work on Alzheimer's disease is concerned with evaluating the changes in the reservoir computing capacity of the network as the network progressively deteriorates dues to synaptic loss. So far, we have examined a simple classification task; in future work, we will gain a more comprehensive understanding of the relationship between synaptic loss and neural computation by extending this work to the battery of tasks described above. In addition, in collaboartion with experimental partners at the University of Freiburg, we will be investigating the dynamic commonalities between Alzheimer's disease and epilepsy, and characterising the changes in computational capacity as the network traverses bifurcation points. With respect to Parkinson's disease, in our previous work we used a genetic algorithm to generate large ensembles of model configurations for healthy and diseased basal ganglia networks (Bahuguna et al., 2017). In our current work, we are applying dimension reduction techniques and graph analysis to determine the simplest manipulations to convert healthy configurations to diseased ones, and vice versa. This will shed light on the path of disease progression, and allow us to propose hypotheses for possible therapeutic interventions. In a follow-up study, we will exploit the interface to robotic simulators to examine the effect of network degeneration on motor behaviour.

Simulation technology

The primary method of investigation for the research topics described above is the simulation of spiking neuronal networks. Therefore, a long term secondary focus of my work is the development of the corresponding HPC amenable simulation technology, NEST (www.nest-simulator.org). NEST is the main large-scale simulator of the EU Flagship "Human Brain Project", and has been used to produce hundreds of research papers in computation neuroscience. My work in this area includes the design of data structures, communication patterns and efficient algorithms for synaptic plasticity and neuronal dynamics (Morrison et al., 2005; Plesser et al., 2007; Potjans et al., 2010; Kunkel et al., 2012). This has enabled us to carry out the world's largest general purpose neuronal network simulations (1.73×10^9 neurons and 1.04×10^{13} synapses), exploiting the full scale of some of the world's largest supercomputers: JUQUEEN at Forschungszentrum Jülich and Japan's K computer (Kunkel et al., 2011; Helias et al., 2012; Kunkel et al., 2014). More recently, we have extended the applicability of NEST by creating interfaces to virtual robotic simulators, enabling closed-loop experiments of agents learning in environments, and to interactive visualization and steering applications (Weidel et al., 2016; Diaz-Pier et al., 2016). We are currently extending this work to allow NEST models to interact with an AI benchmarking framework (Open AI Gym).

Through many years of using, and supporting the use of, NEST for research and teaching, it emerged that a common problem of neuroscience researchers was the difficulty in creating new neuron and synapse models for NEST, as this required at least reasonable C++ skills. To address this issue, in collaboration with Prof. Rumpe and initially financed by a JARA seed fund, we have been developing a domain specific language that allows a user to express neuron models in a succinct and intuitive fashion; the C++ code, including the approprate numerical solver for the specified dynamics, is then generated automatically (Plotnikov et al., 2016; Blundell et al., 2018). The next target for this research project is models of synaptic plasticity.

Neuromorphic Computing

Due to my research in both computational neuroscience and simulation technology, in recent years I have become increasingly interested in neuromorphic computing. This holds out the promise of enabling “cognitive” computing - i.e. employing creativity and generalisation - using low amounts of energy. My current and future contributions to this exciting new field fall into four distinct areas. Firstly, I am involved with requirements and use case definition. Often the hardware developers lack the neuroscience background to ascertain what are the crucial aspects of neurobiology, being arbitrarily complex on all scales, that should be primary goals for an innovative neuromorphic hardware system. As a computation neuorscientist, I am in a good position to mediate between ’wet’ neuroscience and hardware design.

A second area of interest is verification and benchmarking, applying my experience in NEST of testing accuracy and evaluating performance, to networks implemented in neuromorphic systems. In particular, cross checking against the same network implemented in NEST has proved a valuable tool for identifying accuracy issues (Trensch et al., 2018).

Thirdly, we will not be able to realise the potential of neuromorphic computing platforms if they can only be operated and programmed by hardware developers; thus attention must be paid to the usability of such systems. In addition to encouraging system designers to have sensible APIs and interfaces to higher level languages like PyNN, we are exploiting the domain specific language NESTML, as described above. The abstraction achieved through the modelling enables us to produce/create/generate efficient representation not only for NEST but also for neuromorphic backends such as SpiNNaker.

Finally, a neuromorphic hardware system requires algorithms that can run on it. We investigate what functional networks can enable nueuromorphic computing, and whether they are compatible with existing or proposed hardware constraints. For example, given that biological levels of connectivity are likely to be a challenge for most hardware platforms, we are assessing whether the performance of a reservoir computing network can be maintained when using fewer, but stronger, synapses, and uncovering the limits of this scaling approach.

TEACHING PROPOSAL FOR RWTH AACHEN

- **Lecture: Neural and neuro-inspired computing**
  Weekly course with accompanying exercises, covering topics such as reservoir computing, supervised and reinforcement learning in neuronal networks, and the links from neuroscience to artificial neural networks
- **Practical: Simulation of biological neuronal networks**
  Block course with accompanying lectures, covering all necessary aspects to specify, run and analyze spiking neuronal network models using the simulator NEST, including interface to virtual robotic systems
- **Seminar: Functional network models**
  Topics covering seminal and recent studies presenting spiking neuronal networks capable of performing tasks such as navigation, classification and pattern recognition. Specific attention will be paid to extracting the underlying mechanisms.
- **Seminar: High-Performance simulation in the Life Sciences**
  Topics covering advances in HPC-amenable algorithms for simulating biological processes
- **(Co-)Supervision of BSc and MSc projects and internships**