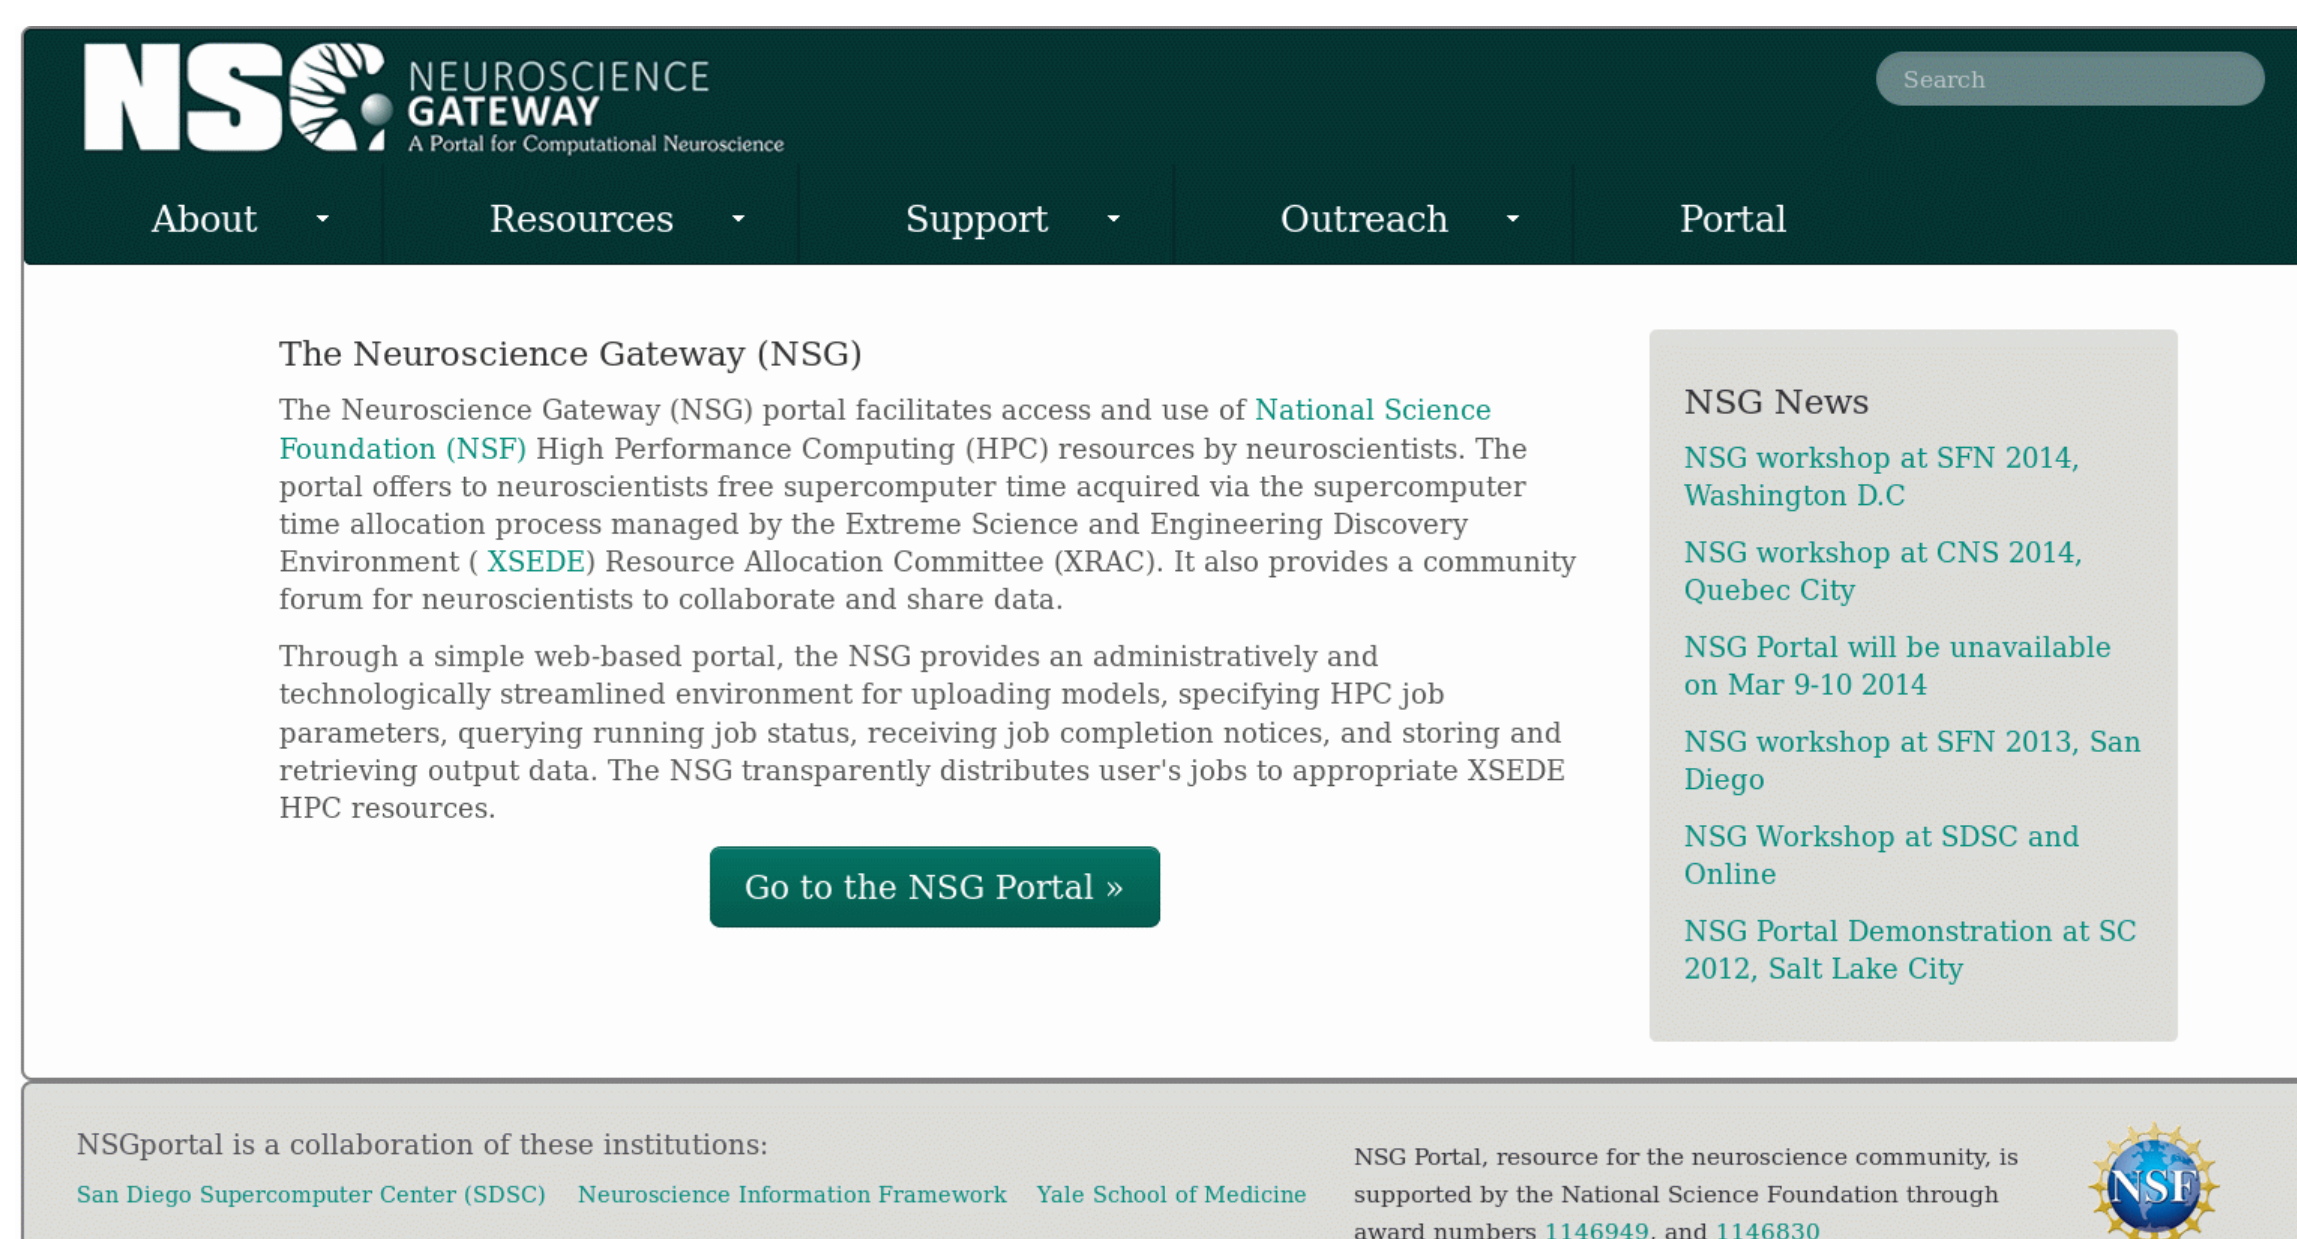


# High Performance Computing in Neuroscience via the Neuroscience Gateway Portal

\*N.T. Carnevale<sup>1</sup>, A. Majumdar<sup>2</sup>, S. Sivagnanam<sup>2</sup>, K. Yoshimoto<sup>2</sup>, V. Astakhov<sup>2</sup>, A. Bandrowski<sup>2</sup>, M. Martone<sup>2</sup>

<sup>1</sup>Yale University School of Medicine, New Haven, CT; <sup>2</sup>UCSD, San Diego, CA

NIF



## Abstract

We are developing the **Neuroscience Gateway (NSG) Portal** <http://www.nsgportal.org/> [NSG] to facilitate the use of high performance computing (HPC) resources by computational neuroscientists. Access to such resources is essential for computationally intensive tasks such as simulating large-scale networks or exploring high-dimensional parameter spaces. However, potential users face a high entry barrier because of numerous administrative and technical details, which include

- writing allocation proposals for computer time on HPC resources
- optimally installing and configuring simulators on parallel hardware
- managing data transfer, storage, and retrieval
- dealing with complex HPC OS and hardware issues
- mastering batch processing environments and authentication policies

To make matters worse, many of these details vary from one site to another.

NSG reduces or eliminates these side-issues by streamlining administrative procedures and providing a browser-based interface that hides technical details. This allows users to focus on scientific issues that are relevant to their research.

We have been obtaining time allocations for the NSG Portal on XSEDE HPC resources, from which each NSG user is given an initial allotment of 20,000 core hours per year, subject to adjustment in light of actual usage by all NSG users. Those who have independently obtained their own HPC allocation may manage that time via the NSG.

Simulators currently installed and available to NSG users include Brian 1.4.1, NEST 2.2.1, NEURON 7.2 and 7.3, PGGENESIS 2.3, and PyNN 0.7.5, and we plan to install MOOSE when its parallel version is available. Through NSG's browser-based interface, users can

- upload input files or model source code as plain text or collections of files in a pkzip archive
- specify code-specific input files
- specify job submission parameters (e.g. number of cores and nodes, expected wall clock time for job completion)
- monitor the status of submitted jobs (users are also notified of job completion by email)
- download job output files as a pkzip archive of the job's working directory, which contains input and output files

We will present performance metrics, usage statistics, and specific examples of productive use of the NSG for neuroscience research. We welcome suggestions for other neuroscience simulators or software packages, and invite suggestions for ways to increase the utility of the NSG in neuroscience research.

## Milestones and usage


12/2012 NSG opened for use.

3/2013 Initial 50k CPU hour allocation consumed.

11/2013 Total CPU usage >350k hours since start.

2014 XSEDE awarded 1.6e6 service units on SDSC's Trestles for calendar year (2.4 GHz AMD Magny-Cours processor, 32 cores and 64 GB memory/node)


11/2014 ~200 users, ~40 "active"  
~120K service units still available on Trestles,  
~61K hours on TACC's Stampede (~2x faster than Trestles)



## Performance

These figures illustrate how the number of processors (NP) used in a parallel simulation affects runtime on a desktop PC and the NSG. These particular results were obtained from simulations with NEURON of two spiking network models in which cell coupling was implemented with spike-triggered events; similar results were obtained with a model of a 2D rectangular mesh of myocytes in which adjacent cells were coupled by rectifying gap junctions.

In theory runtime will vary as 1/NP if interprocessor communication overhead is negligible compared to computation time (solid and dashed black lines in these figures). In practice, cache effects often result in supralinear speedup for "small" values of NP, as seen here--note that actual NSG run times (blue lines) are shorter than predicted by theory (black dashed lines).



## How NSG helps advance neuroscience

The Neuroscience Gateway can benefit the broader neuroscience research community in several ways, e.g.:

- Simple interface provides streamlined access to HPC resources, so investigators can focus on their research.
- Can be used by investigators with limited local (university-level) resources to address questions that require access to large scale, advanced systems.
- Can be used by simulator developers to test/debug, benchmark, and scale codes on large scale resources and, when satisfied, make them available to the broader user community.
- Promotes use of HPC resources in classes, workshops, and tutorials by relieving instructors the need to set up simulators on HPC resources or create new accounts for students/participants.

## Publications enabled by NSG

T. Rumbell, D. Draguljić, J.I. Luebke, P.R. Hof, C.M. Weaver. Compartmental model optimization predicts altered channel densities and kinetics in aged versus young pyramidal neurons of rhesus monkey prefrontal cortex. Poster 372.22/VV86, Monday 11/17/2014 8AM-Noon, Society for Neuroscience Annual Meeting, Washington D.C., Nov. 15-19, 2014.

T. Rumbell, D. Draguljić, J. Luebke, P. Hof, C.M. Weaver. Automatic fitness function selection for compartment model optimization. BMC Neuroscience 15, Suppl 1 (2014): O5.

S. Lee, I. Marchionni, M.J. Bezaire, N. Danielson, M. Lovett-Barron, A. Losonczy, I. Soltesz. GABAergic basket cells differentiate among hippocampal pyramidal cells. Neuron 82:1129-1144, 2014.

L. Ingber, M. Pappalepore, R. Stesiak. Electroencephalographic field influence on calcium momentum waves. Journal of Theoretical Biology 343:138-153, 2014.

## Publications about NSG

T. Carnevale, A. Majumdar, S. Sivagnanam, K. Yoshimoto, V. Astakhov, A. Bandrowski, M. Martone. High performance computing in neuroscience via the Neuroscience Gateway Portal. Poster 187.01/TT50, Sunday 11/16/2014 8AM-Noon, Society for Neuroscience Annual Meeting, Washington D.C, Nov. 15 - 19, 2014.

T. Carnevale, A. Majumdar, S. Sivagnanam, K. Yoshimoto, V. Astakhov, A. Bandrowski, M. Martone. The Neuroscience Gateway Portal--high performance computing made easy. Poster, Computational Neuroscience (CNS) 2014 Annual Meeting, Quebec City, Canada, July 26-31, 2014.

S. Sivagnanam, A. Majumdar, K. Yoshimoto, V. Astakhov, A. Bandrowski, M. Martone, N.T. Carnevale. Early experiences in developing and managing the neuroscience gateway. Journal of Concurrency and Computation: Practice and Experience, May 2014, UR: <http://dx.doi.org/10.1002/cpe.3283>.

N.T. Carnevale, A. Majumdar, S. Sivagnanam, K. Yoshimoto, V. Astakhov, A. Bandrowski, M. Martone. The Neuroscience Gateway Portal: facilitating access to high performance computing resources. Poster, Society for Neuroscience Annual Meeting, San Diego, Nov. 9-13, 2013.

S. Sivagnanam, K. Yoshimoto, A. Majumdar, N. T. Carnevale, V. Astakhov, M. Martone, A. Bandrowski. A Neuroscience Gateway: software and implementation. Proceedings XSEDE13 Gateway to Discovery, San Diego, CA, July 22-25, 2013.

S. Sivagnanam, A. Majumdar, K. Yoshimoto, N. T. Carnevale, V. Astakhov, A. Bandrowski, M. Martone. Introducing The Neuroscience Gateway. Proceedings International Workshop on Science Gateways, Zurich, Switzerland, June 3-5, 2013.

N.T. Carnevale, S. Sivagnanam, K. K. Yoshimoto, V. Astakhov, A. E. Bandrowski, M. E. Martone, A. Majumdar. A Neuroscience Gateway for high performance computing. Poster, Society for Neuroscience Annual Meeting, New Orleans, October 13-17, 2012.

## Work in progress

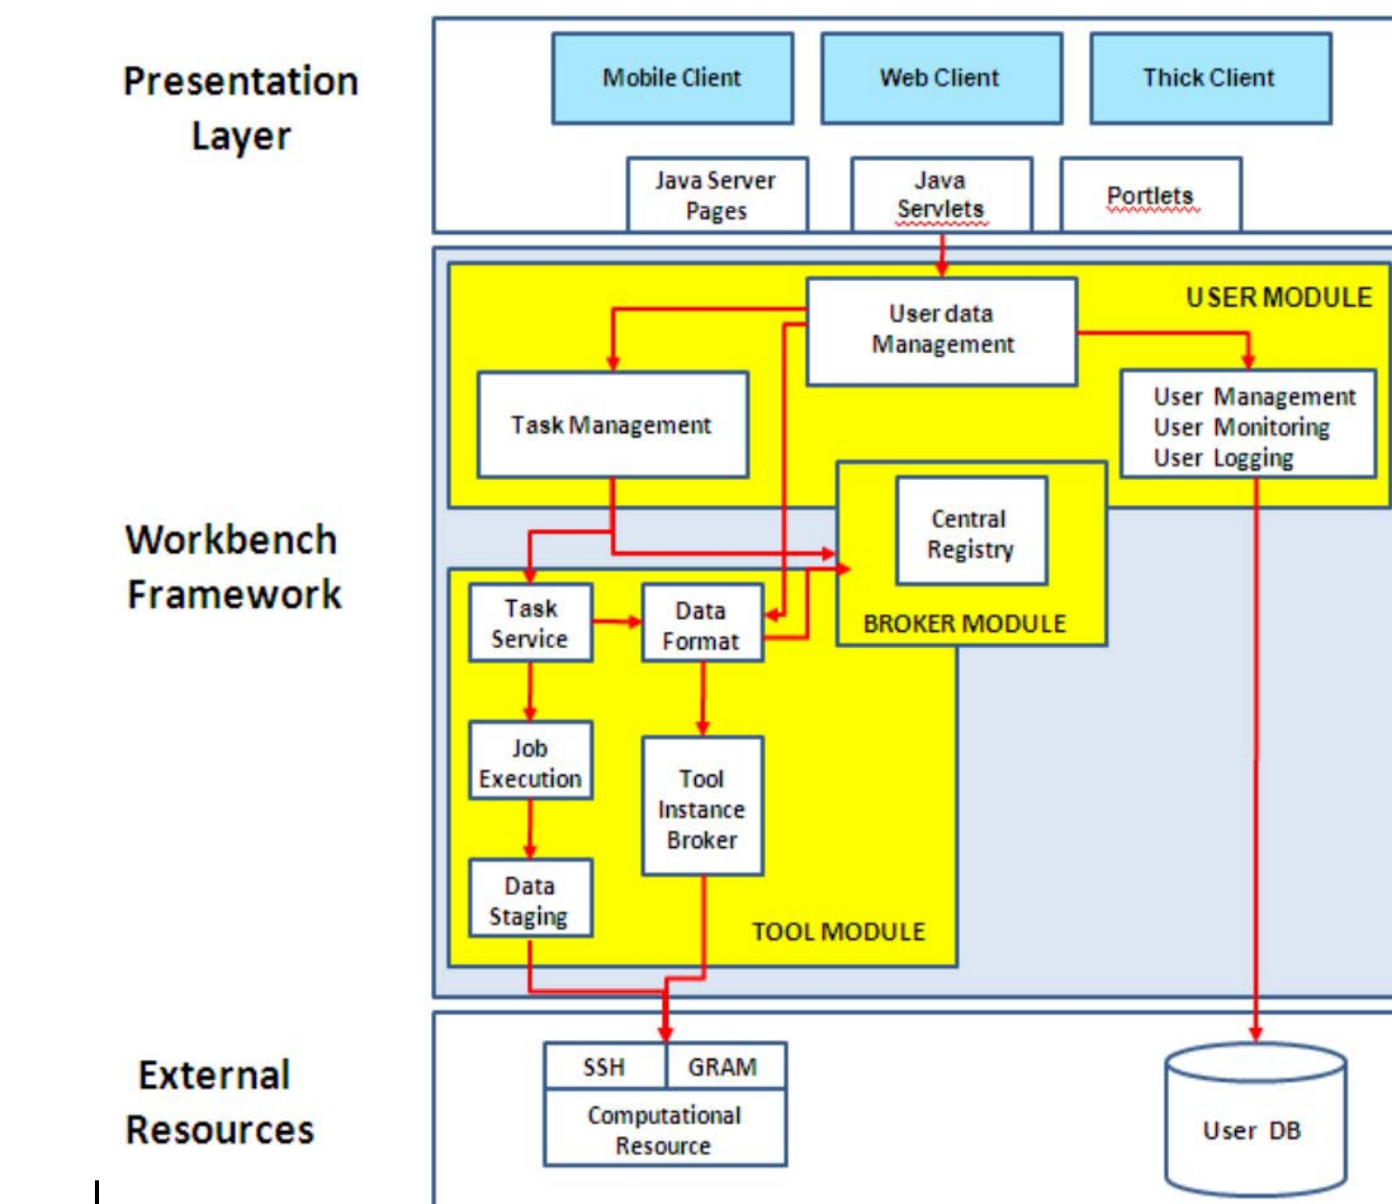
- debugging capability
- sharing of output results, e.g. with NIF
- installation of other simulators e.g. MOOSE
- other enhancements driven by user suggestions and feedback

## Technical details

NSG hides C/HPC-related complexities within the black box of a gateway architecture that is based on and extends the CIPRES Science Gateway framework [CIPRES, Miller 2010]. CIPRES was developed at the San Diego Supercomputer Center (SDSC) for the computational phylogenetics research community. All enhancements and modifications implemented for NSG are contributed back to CIPRES for future use.

Some features of NSG differ from CIPRES, such as: any NSG model may involve multiple files and a complex directory hierarchy; some simulators require compilation of model source code, e.g. NEURON's mod files; results are automatically stored in SDSC's cloud storage; user files are automatically deleted based on duration of inactivity.

NSG's architecture uses WF, which contains modules to manage submission of jobs to analytical tools on computational resources, and modules to manage queries on data resources.



Presentation Layer: Supports access by browser clients.

User Module: Manages all user-initiated activities, passing queries and tasks to executive portions of the infrastructure; stores user and task information; manages user roles, account assignments, data sharing between accounts, access to tools and data sources (which may be proprietary).

Tool Module: Translates user-submitted tasks into command lines, submits commands and user data to appropriate compute engine for execution.

External Resources: Can be accessed via SSH, GRAM/Globus [GLOBUS], etc..

## Acknowledgements and References

NSF collaborative grant

DBI 1146949 PI A. Majumdar, SDSC; Co-PI M. Martone, NIF, UCSD

DBI 1146830 PI T. Carnevale, Yale University School of Medicine

T. Carnevale was also partly supported by NIDCD DC009977 PI G. Shepherd

SDSC Internal grant to S. Sivagnanam and K. Yoshimoto for cloud storage  
Technical support for NEURON via NINDS NS11613 PI M. Hines and the Blue Brain Project

[CIPRES] <http://www.phylo.org/>

[GENESIS3] <http://www.genesis-sim.org/GENESIS/G3/>

[GLOBUS] <http://globus.org/>

[GRAM] <http://dev.globus.org/wiki/GRAM>

[Miller 2010] M. Miller, W. Pfeiffer, T. Schwartz, "Creating the CIPRES Science Gateway for inference of large phylogenetic trees," Gateways Computing Environments Workshop (GCE), 2010, pp 1-8, New Orleans, LA, 14 Nov., 2010. DOI 10.1109/GCE.2010.5676129.

[ModelDB] <http://senselab.med.yale.edu/modeldb/>

[MOOSE] <http://moose.sourceforge.net/>

[NEST] [http://www.nest-initiative.unifreiburg.de/index.php/Software>About\\_NEST](http://www.nest-initiative.unifreiburg.de/index.php/Software>About_NEST)

[neuroConstruct] <http://www.neuroconstruct.org/>

[NEURON] <http://www.neuron.yale.edu/>

[NIF] <http://neuinfo.org/>

[NSG] S. Sivagnanam, A. Majumdar, K. Yoshimoto, V. Astakhov, A. Bandrowski, M. E. Martone, and N. T. Carnevale. Introducing the Neuroscience Gateway, IWSG, volume 993 of CEUR Workshop Proceedings, CEUR-WS.org, 2013

[PyNN] <http://neuralensemble.org/trac/PyNN/>

[WBC] <http://wholebraincatalog.org/>

[WF] <http://www.ngbw.org/wbframework/>

[XSEDE] <http://www.xsede.org/>

