SEARCHLIGHT GOES GPU - FAST MULTI-VOXEL PATTERN ANALYSIS OF FMRI DATA

Anders Eklund[1], Malin Björnsdotter[2,3], Johannes Stelzer[4], and Stephen LaConte[1,5]

[1]*Virginia Tech Carilion Research Institute, Roanoke, Virginia, United States,* [2]*University of Gothenburg, Göteborg, Sweden,* [3]*Nanyang Technological University, Singapore, Singapore,* [4]*Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig, Germany,* [5]*School of Biomedical Engineering & Sciences, Virginia Tech-Wake Forest University, Blacksburg, Virginia, United States*

INTRODUCTION Graphics processing units (GPUs) are used today in a wide range of applications, mainly because they can dramatically accelerate parallel computing, are affordable and power efficient. In the field of medical imaging, GPUs can make computationally demanding algorithms practicle, e.g. for image reconstruction[1]. For fMRI, GPUs have recently been used to speedup non-parametric statistical methods[2,3], which can be more reliable than parametric methods used by most software packages[4]. Here we demonstrate the utility of GPUs for multivariate fMRI analysis. The searchlight algorithm[5] is a popular choice for locally-multivariate decoding of fMRI data. For each voxel, a classifier is trained to discriminate between different brain states, by using voxels within a neighborhood search volume. A performance measure, typically the classification accuracy, is then saved in the center voxel. A substantial drawback of the searchlight is that it is computationally demanding. This is especially true for large searchlight spheres, non-linear classifiers, cross validation schemes and statistical permutation testing. Here we therefore present a GPU implementation of the searchlight algorithm, which is over 8000 times faster than a simple Matlab implementation and about 21 times faster than a parallelized Matlab implementation.

METHODS Voxel-wise p-values for the searchlight can be calculated from the classification accuracy, through the binomial distribution. This can, however, be overly permissive if cross-validation is applied[6]. Another problem is in how to correct the p-values, both for multiple comparisons and for cluster threshoding. A non-parametric approach, such as a random permutation test, can solve these problems empirically, by deriving a null distribution from the data itself[2]. However, p-values that are corrected for multiple comparisons require several thousand permutations. According to recent work[6], 10,000 permutations containing training and evaluation of a support vector machine (SVM) classifier would take about 277 hours on a CPU.

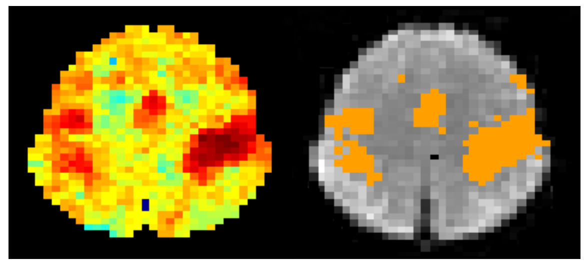

Fig 1. **Left:** A single searchlight computation, un-thresholded due to the lack of p-value estimates and correction for multiple comparisons. **Right:** The information map thresholded at a cluster level of p = 0.05, corrected for multiple comparisons using 10,000 permutations.

In each brain voxel, a linear one layer artificial neural network (ANN) was trained with a searchlight (diameter 3 voxels, total of 19 voxels) to classify the brain activity. Leave-one-out cross validation was applied to estimate the classification accuracy. The GPU implementation was done with the CUDA programming language, such that each GPU thread performed the calculations for one voxel. fMRI data of a simple hand motor task (20 s rest, 20 s activity, four periods), collected with a Philips Achieva 1.5 T MR scanner, was used to test our implementations. The size of the dataset was 80 volumes of a resolution 64 x 64 x 22 voxels. The physical size of each voxel was 3.75 x 3.75 x 3.75 mm and the repetition time was 2 seconds. Prior to the decoding, motion correction and cubic detrending was applied. Whitening with a voxel specific AR(4) model was applied prior to the permutations (as a permutation test requires that the labels can be exchanged under the null hypothesis). The GPU implementation was compared to a straightforward Matlab implementation and a parallelized Matlab implementation using the open multiprocessing (OpenMP) library, which makes it possible to use all the CPU cores. All processing was done using a workstation equipped with a 3.5 GHz Intel Core i7 CPU and an Nvidia GTX 680 graphics card. For the multi-GPU implementation, two Nvidia GTX 690 graphics cards were used (yielding a total of 6144 processor cores in a single PC).

RESULTS The resulting processing times for the different implementations are given in Table 1. The information map thresholded at p = 0.05, corrected for multiple comparisons, is given in Fig. 1 and the estimated maximum null distribution of the classification accuracy is given in Fig 2.

Table 1	**Matlab**	**Matlab OpenMP**	**Single GPU**	**Multi-GPU**
Searchlight ANN	**962 s**	**2.5 s**	**0.11 s**	**Not performed**
Searchlight ANN, 10 000 permutations	**111 days**	**7 h**	**19 min**	**5 min**

DISCUSSION We have presented a GPU implementation of the searchlight algorithm. The GPU implementation is 8300 times faster than a straightforward Matlab implementation and about 21 times faster than a parallelized Matlab implementation. By applying a random permutation test with 10,000 permutations, the maximum null distribution was estimated to calculate corrected p-values (both for voxels and clusters). Here we have used an ANN due to its simplicity, in future work we will also implement GPU based searchlight for SVM[7].

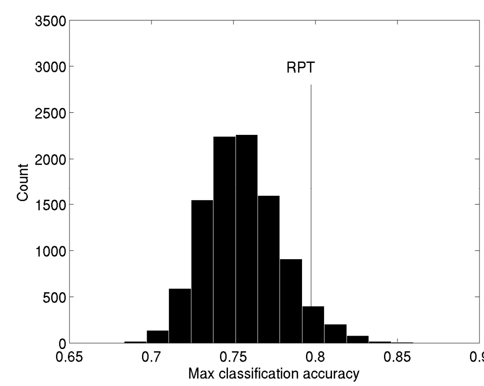

Fig 2. The estimated null distribution of the maximum cross-validated classification accuracy and the voxel-wise threshold for corrected p = 0.05 (marked RPT).

REFERENCES 1 Stone, S., Haldar, J. Tsao S., Hwu, W., Sutton B., Liang Z., Accelerating advanced MRI reconstructions on GPUs, Journal of parallel and distributed computing, 68, 1307-1318, 2008 **2** Nichols, T.E., Holmes, A.P., Nonparametric permutation tests for functional neuroimaging: a primer with examples, Human Brain Mapping, 15, 1-25, 2002 **3** Eklund, A., Andersson, M., Knutsson, H., fMRI analysis on the GPU – Possibilities and challenges, Computer methods and programs in biomedicine, 105, 145-161, 2012 **4** Eklund A., Andersson, M., Josephson, C., Johannesson, M., Knutsson H., Does parametric fMRI analysis with SPM yield valid results? – An empirical study of 1484 rest datasets, NeuroImage, 61, 565-578, 2012 **5** Kriegeskorte, N., Goebel, R., Bandettini, P. Information-based functional brain mapping, PNAS, 103, 3863-3868, 2006 **6** Stelzer, J., Chen, Y., Turner, R., Statistical inference and multiple testing correction in classification-based multi-voxel pattern analysis (MVPA): Random permutations and cluster size control, NeuroImage, doi: 10.1016/j.neuroimage.2012.09.063 **7** LaConte, S., Strother, S., Cherkassky, V:, Anderson, J., Hu, X. Support vector machines for temporal classification of block design fMRI data, NeuroImage, 26, 317-329, 2005