

Spatial and Temporal Feature Extraction for Brain Decoding using CUDA

Itir Onal¹, Alptekin Temizel², Fatos T. Yarman Vural¹

¹Department of Computer Engineering, Middle East Technical University, TURKEY

²Graduate School of Informatics, Middle East Technical University, TURKEY




Motivation

Brain decoding is the process of predicting cognitive states from medical data (fMRI, EEG, etc.) while the subjects are presented with stimulus (picture, audio, etc.). A typical fMRI experiment consists of thousands of voxels and hundreds of samples. Solving regression for all samples of all voxels serially requires huge amount of time.

Spatial and Temporal Features

BOLD response from a seed voxel $\bar{r}(s_i, \bar{l}_j)$ is represented as a linear combination of its nearest neighbors. Arc weights $a_{i,j,k}$ represent both spatial and temporal relationships and they are estimated by solving linear regression.



Solve the equation for all samples and all voxels:

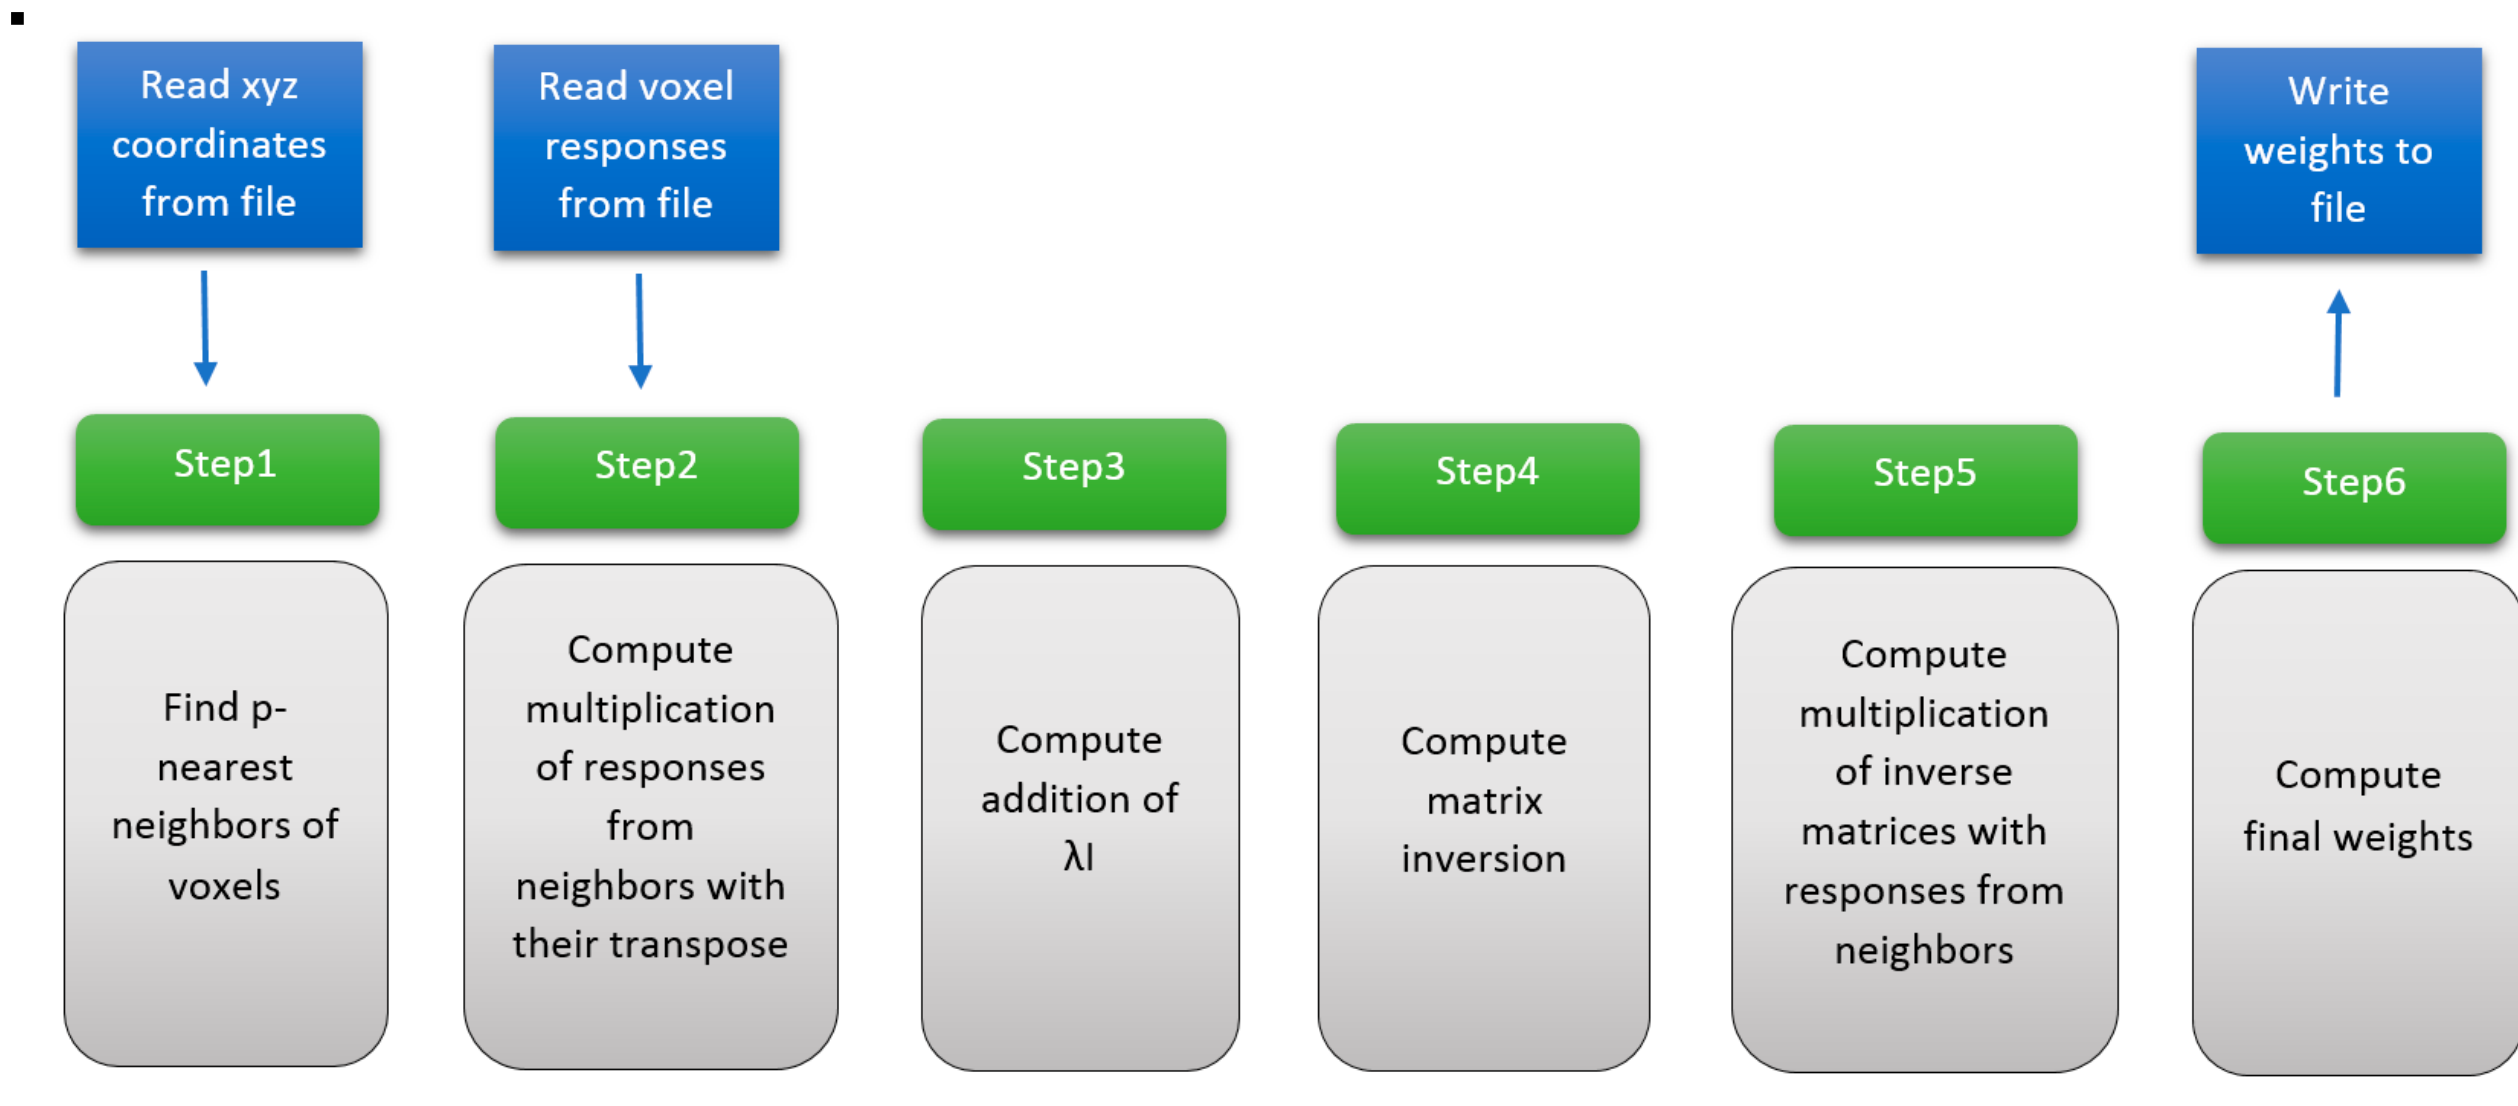
$$\bar{r}(s_i, \bar{l}_j) = \sum_{\bar{l}_k \in \eta_p} a_{i,j,k} \bar{r}(s_i, \bar{l}_k) + \bar{\varepsilon}_{i,j}$$

Closed form solution of ridge regression:

$$\bar{a}_{i,j} = (R_{i,j}^T R_{i,j} + \lambda I)^{-1} R_{i,j}^T \bar{r}(s_i, \bar{l}_j)$$

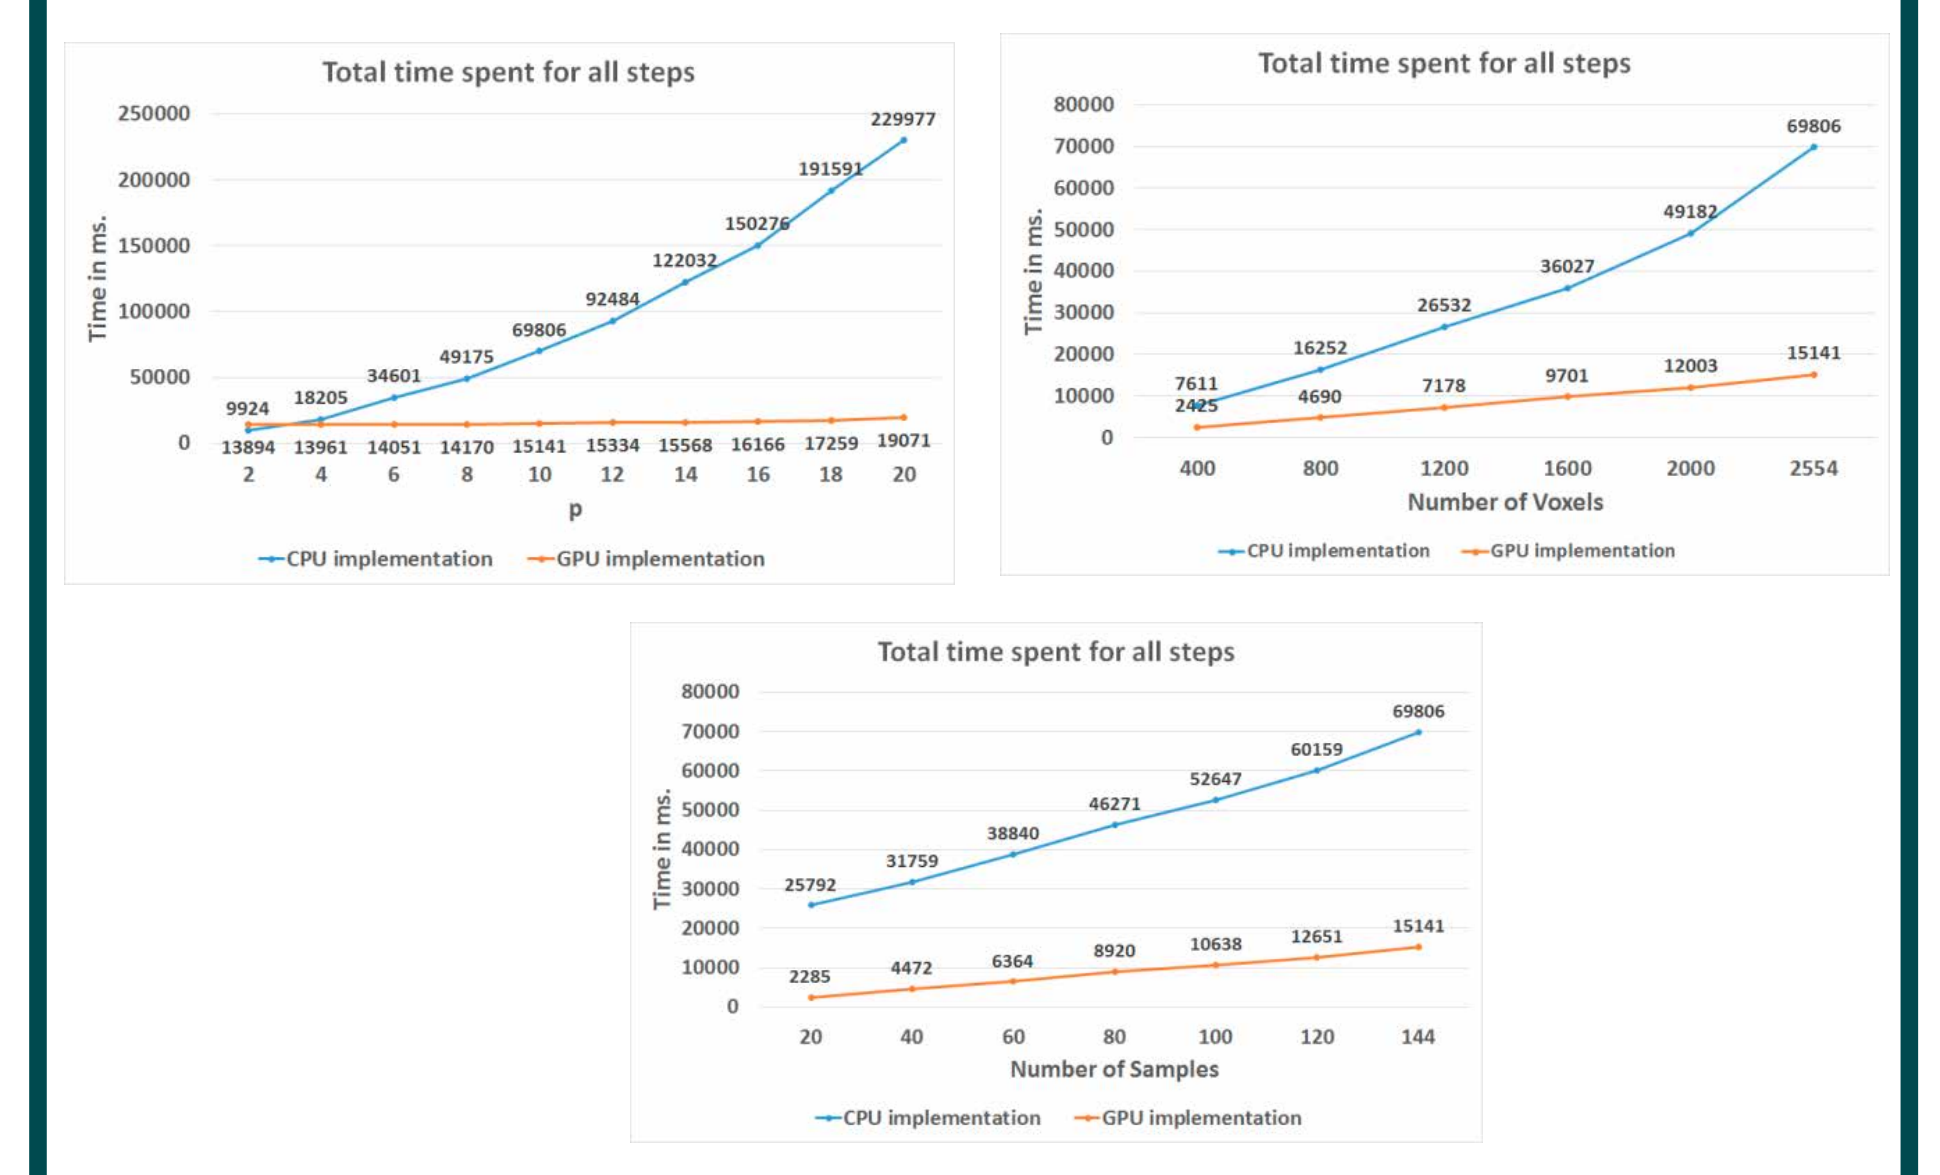
Spatial and Temporal Feature Extraction with CUDA

Flow of algorithm:



Step	Description	Number of neighbors vs. time:	Number of voxels vs. time:	Number of samples vs. time:
Step1	M threads are employed. 1. Compute $M \times M$ Euclidean distances matrix (each thread computes a single row) 2. Find p-nearest neighbors of all voxels (each thread finds p-nearest neighbors of a single voxel)			
Step2 + Step3	$M \times N \times p$ threads are employed 1. Compute $R_{i,j}^T R_{i,j}$ (each thread obtains a single neighbor data and multiplies it with $R_{i,j}$) 2. Compute $R_{i,j}^T R_{i,j} + \lambda I$ (each thread adds a single element to a diagonal in a row)			
Step4	CUBLAS functions are used 1. Compute LU factorization of $M \times N$ many matrices using cublasSgetrfBatched 2. Compute inversion of $M \times N$ many matrices using cublasSgetriBatched			
Step5 + Step6	$M \times N \times p$ threads are employed 1. Multiply $(R_{i,j}^T R_{i,j} + \lambda I)^{-1}$ and $R_{i,j}^T$ (each thread multiplies a row of former with the matrix $R_{i,j}^T$) 2. Obtain final weights, multiply $(R_{i,j}^T R_{i,j} + \lambda I)^{-1} R_{i,j}^T$ with $\bar{r}(s_i, \bar{l}_j)$			

Results and Discussion



- Total time spent for all steps significantly increases as the number of neighbors and voxels increase
- Number of samples does not affect the performance of Step1
- Number of neighbors (p) does not significantly affect the performance of LU factorization and matrix inversion of CUBLAS library functions
- In all experiments, we obtain a speedup with GPU implementation over CPU implementation

CPU: Intel i7 3770K CPU @3.50 GHz with 32GB memory
GPU: GeForce GTX 670 device with CUDA Runtime Version 6.5

Conclusion

- Extraction of spatial and temporal features relies on solving a ridge regression for different parts of data.
- Parallel implementation with CUDA significantly reduces the time spend for extraction.

References

[1] A. Eklund, M. Bjornsdotter, J. Stelzer and S. LaConte, 'Searchlight goes gpu - fast multi-voxel pattern analysis of fMRI data', *International society for magnetic resonance in medicine (ISMRM)*
[2] O. Firat, I. Onal, E. Aksan, B. Velioglu, I. Oztekin and F. T. Yarman Vural, 'Large scale functional connectivity for brain decoding' *BioMed* 2014
[3] M. B. Aberg and J. Wessberg, 'An evolutionart approach to the identification of informative voxel clusters for brain state discrimination', *IEEE Journal of selected topics in signal processing*, vol. 2, pp 919 – 928, 2008