CENTRE DE RECHERCHE COMMUN
INRIA MICROSOFT RESEARCH

Inria
INVENTEURS DU MONDE NUMÉRIQUE

A-Brain:
Large-scale Joint Genetic and Neuroimaging Data Analysis on Azure Clouds

Project PIs: Gabriel Antoniu, Bertrand Thirion
Contributors: Alexandru Costan, Benoit Da Mota, Radu Tudoran and the Microsoft Azure team from EMIC

Final Meeting, MSR-Inria Centre
8 November 2013

The A-Brain Project: Data-Intensive Processing on Microsoft Azure Clouds

Application

- Large-scale joint genetic and neuroimaging data analysis

Goals

- Application: assess and understand the variability between individuals
- Infrastructure: assess the potential benefits of Azure

Approach

- Optimized data processing on Microsoft's Azure clouds

Inria teams involved

- KerData (Rennes)
- Parietal(Saclay)

Framework

- Joint MSR-Inria Research Center
- MS involvement: Azure teams, EMIC

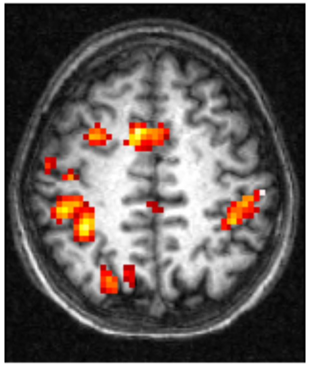

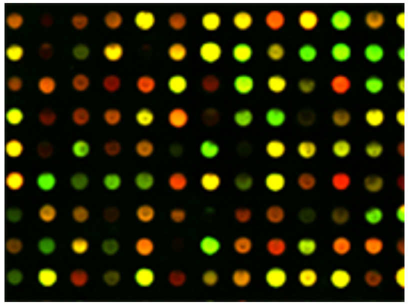

The Imaging Genetics Challenge: Comparing Heterogeneous Information

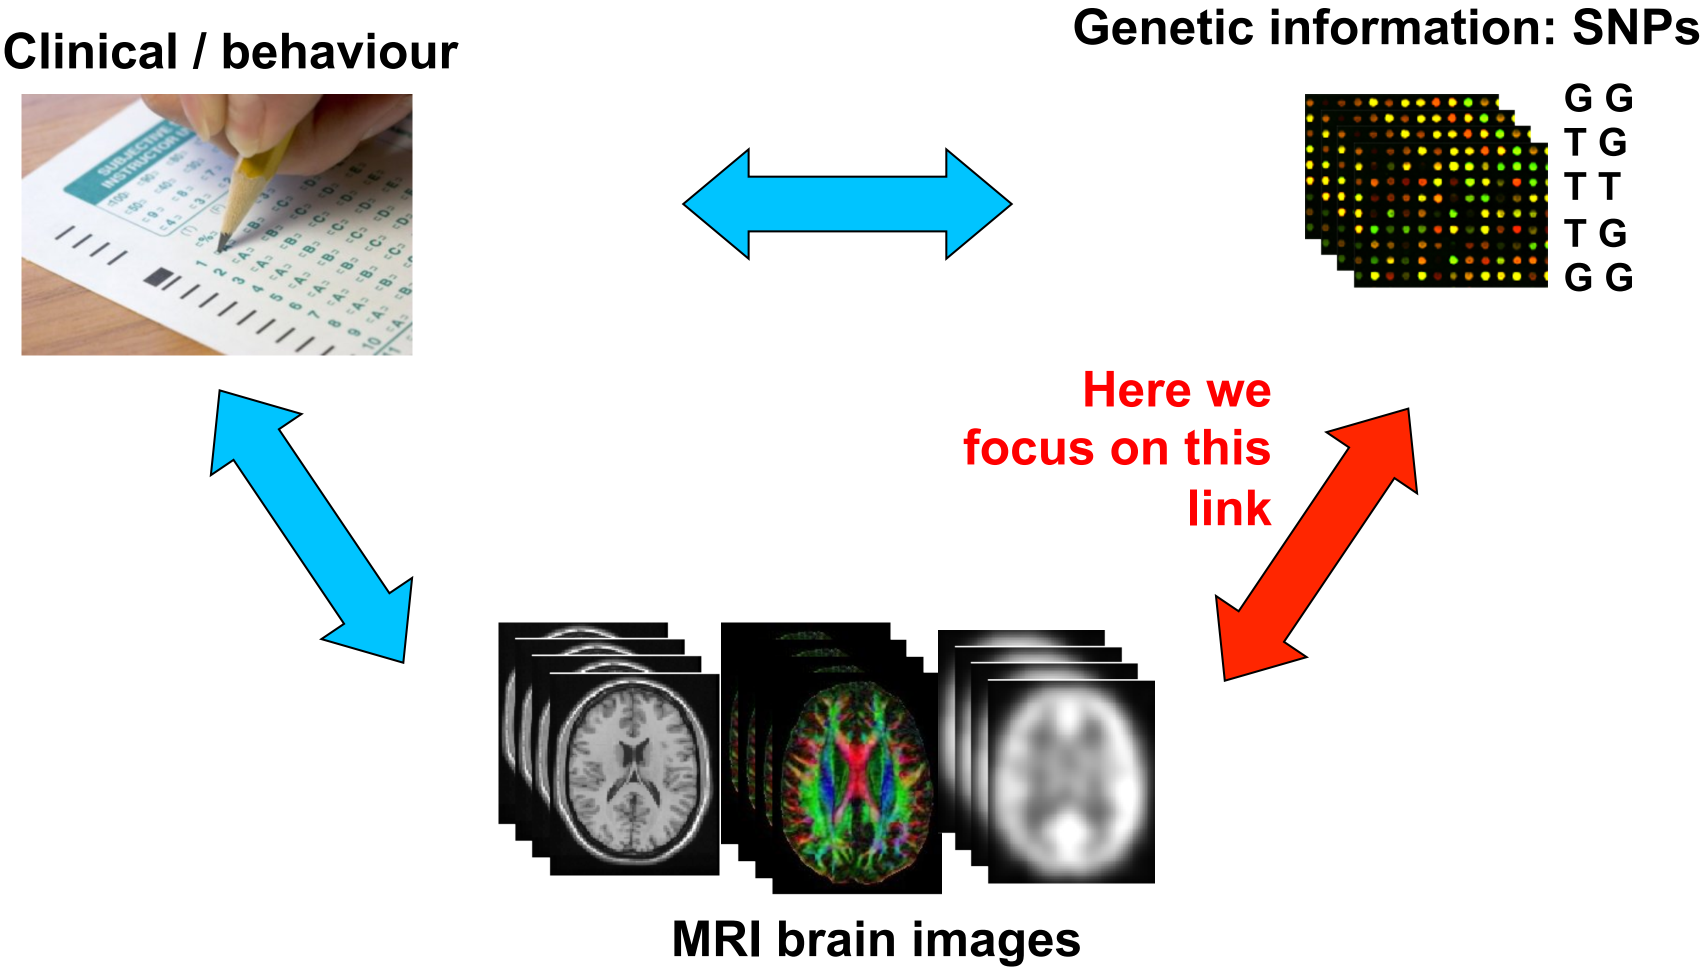

Clinical / behaviour

Genetic information: SNPs

G G
T G
T T
T G
G G

Here we focus on this link

MRI brain images

Neuroimaging-genetics: The Problem

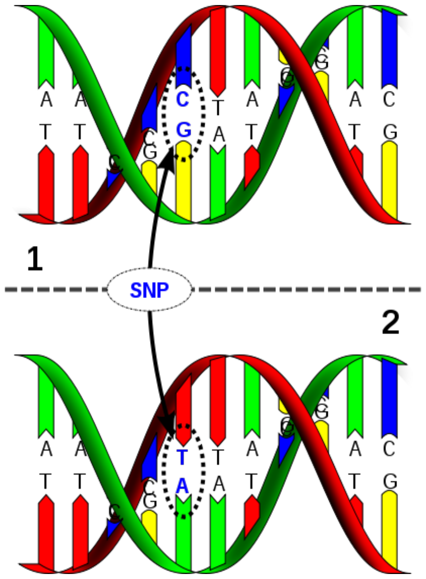

- Several brain diseases have a genetic origin, or their occurrence/ severity related to genetic factors
- Genetics important to understand & predict response to treatment

- Genetic variability captured in DNA micro-array data

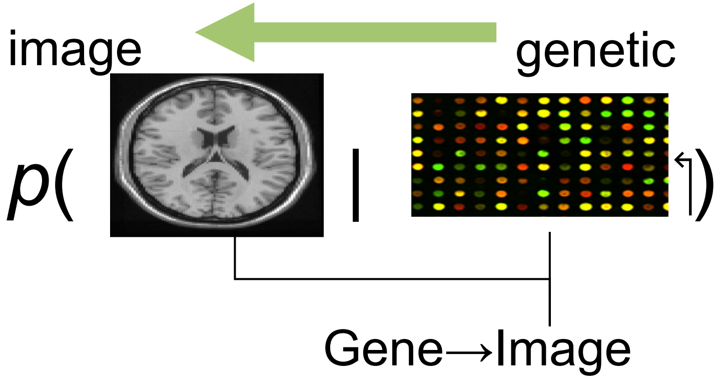

Neuroimaging-genetics studies

- Objective: Find correlation between brain markers and genetic data to understand the behavioral variability and diseases
- Setting: Data pipeline, data organization

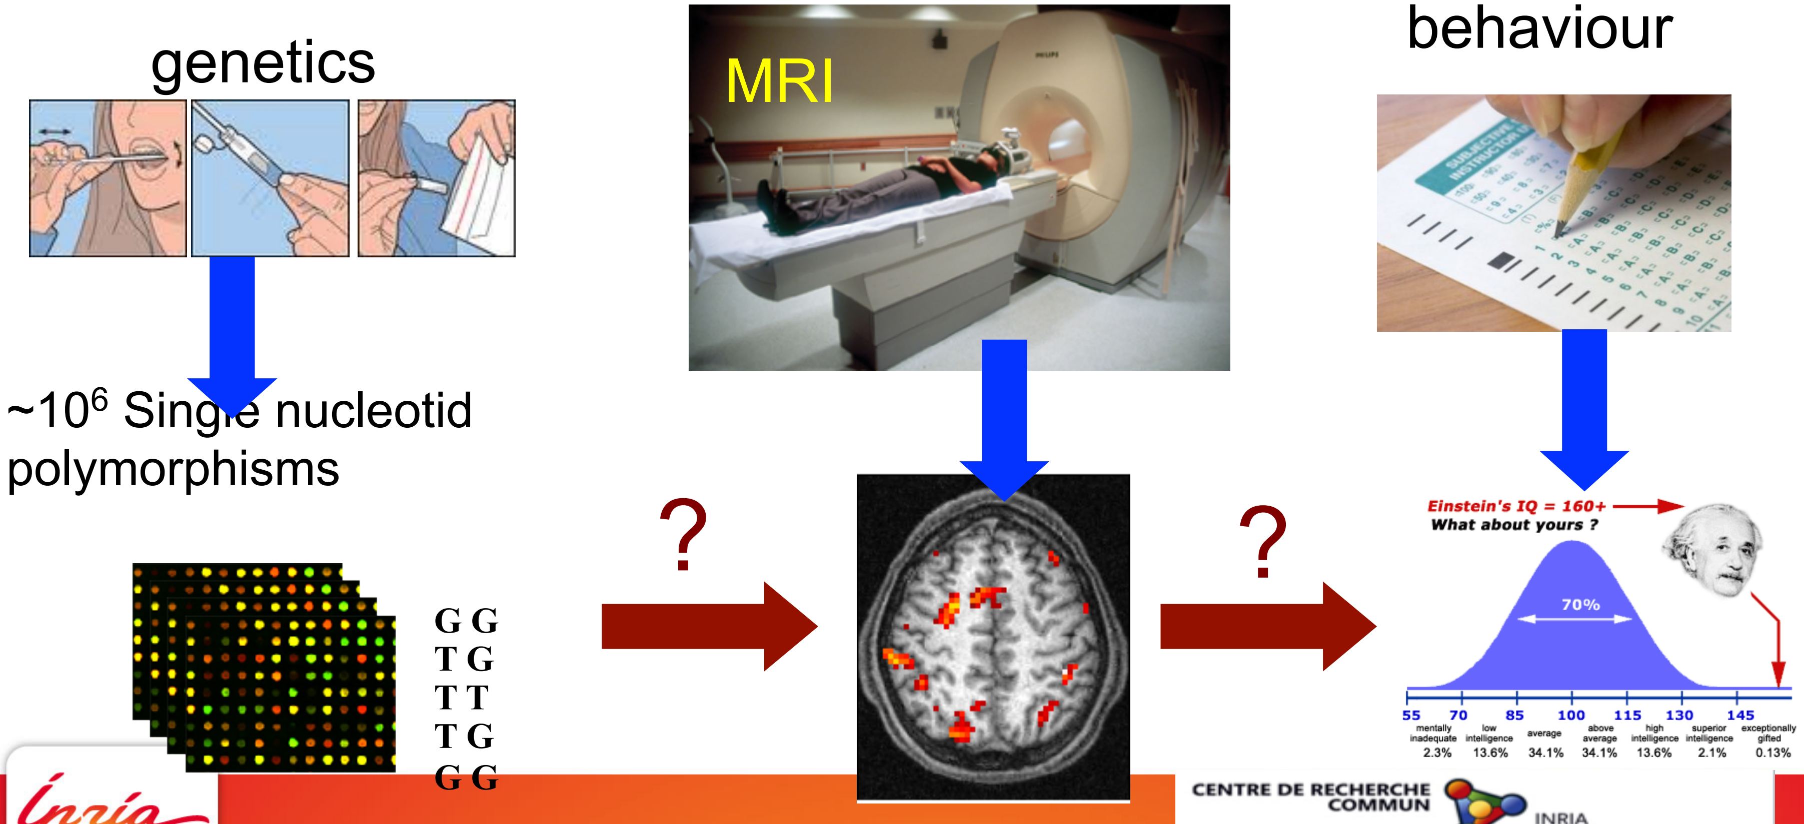

genetics

behaviour

~10^6 Single nucleotid polymorphisms

G G
T G
T T
T G
G G

Statistical analysis for large-scale neuroimaging-genetics

- Image data → 4D to 2D, dimension $n_{voxels} \times n_{subjects}$
- Genetic data → dimension $n_{snps} \times n_{subjects}$
- Statistical question

$n_{voxels} = 10^5$
$n_{snps} = 10^6$
$n_{subjects} = 10^3$

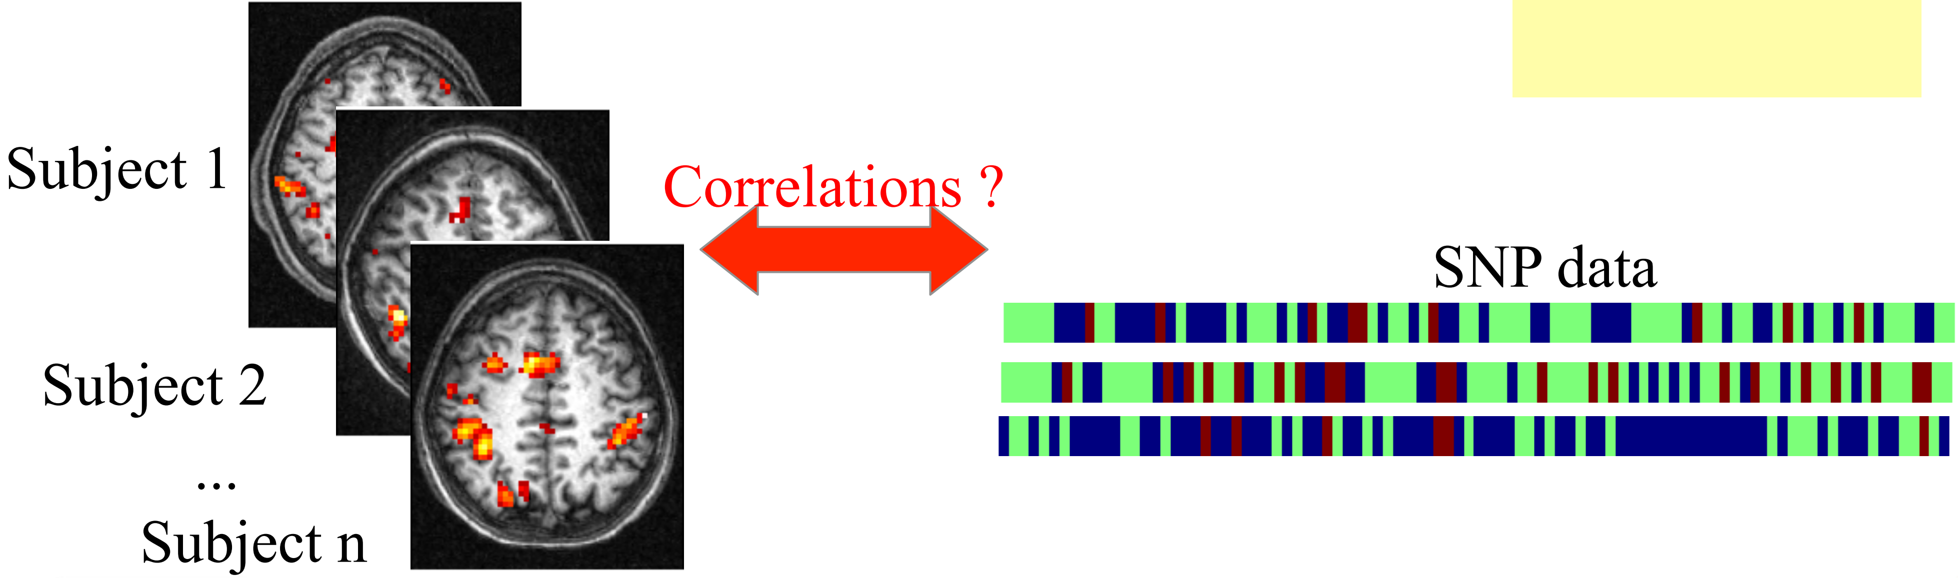

Subject 1

Subject 2

…

Subject n

Correlations ?

SNP data

Approach: A-Brain as Map-Reduce Processing

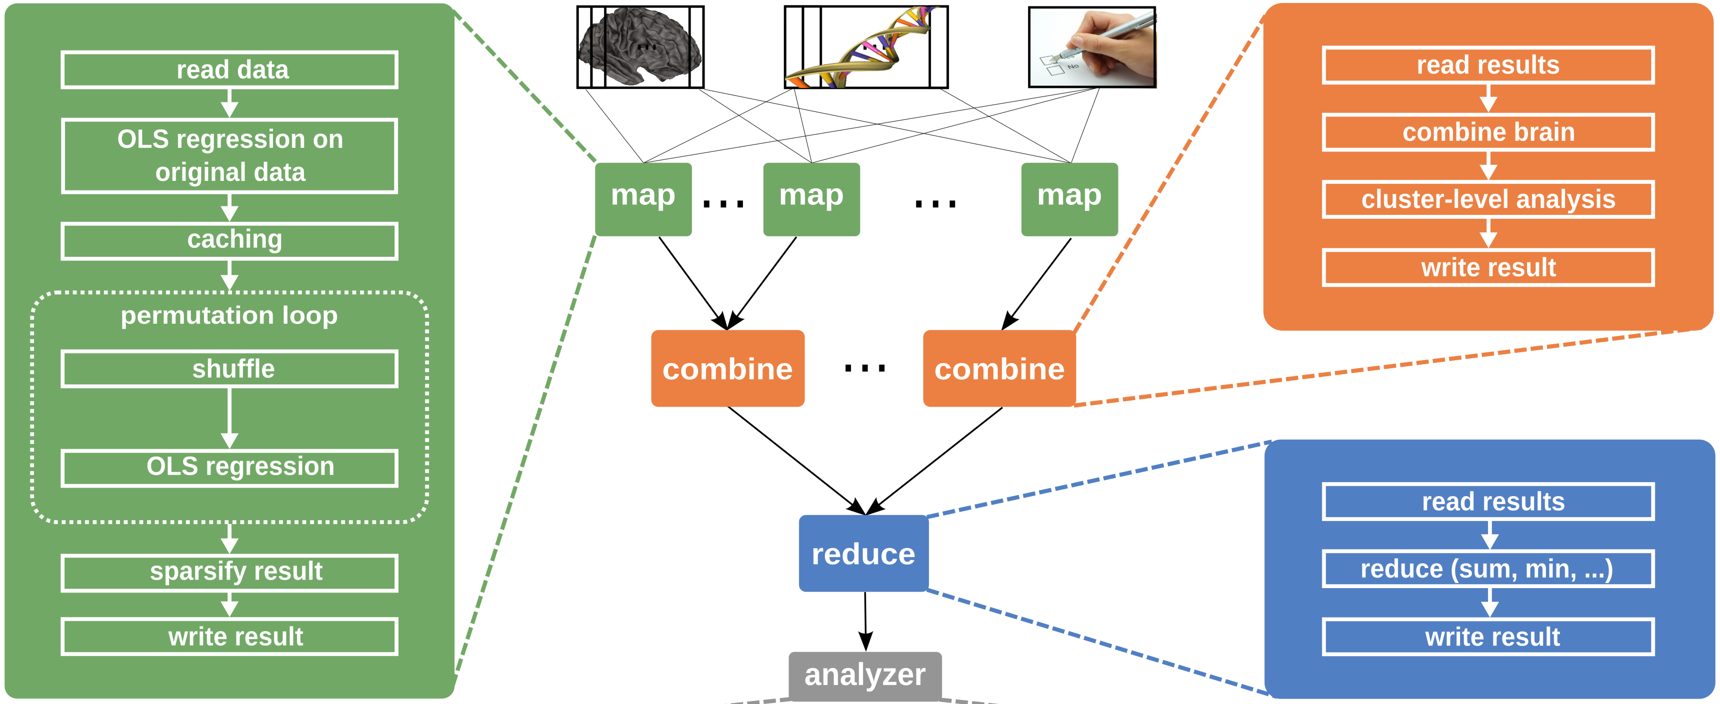

read data
OLS regression on original data
caching
permutation loop
shuffle
OLS regression
sparsify result
write result

read results
combine brain
cluster-level analysis
write result

read results
reduce (sum, min, ...)
write result

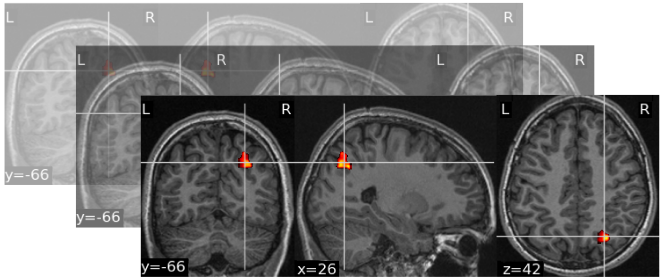

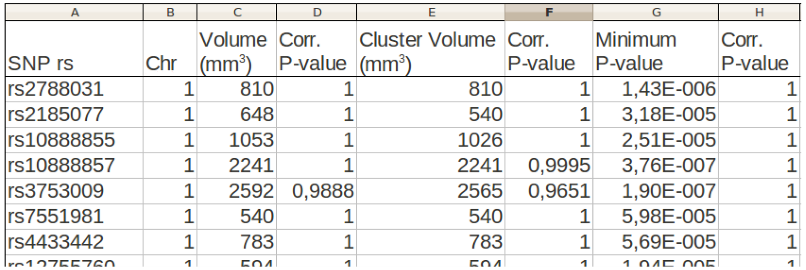

A	B	C	D	E	F	G	H
SNP rs	Chr	Volume (mm³)	Corr. P-value	Cluster Volume (mm³)	Corr. P-value	Minimum P-value	Corr. P-value
rs2788031	1	810	1	810	1	1,43E-006	1
rs2185077	1	648	1	540	1	3,18E-005	1
rs10888855	1	1053	1	1026	1	2,51E-005	1
rs10888857	1	2241	1	2241	0,9995	3,76E-007	1
rs3753009	1	2592	0,9888	2565	0,9651	1,90E-007	1
rs7551981	1	540	1	540	1	5,98E-005	1
rs4433442	1	783	1	783	1	5,69E-005	1

A-Brain as Map-Reduce Data Processing

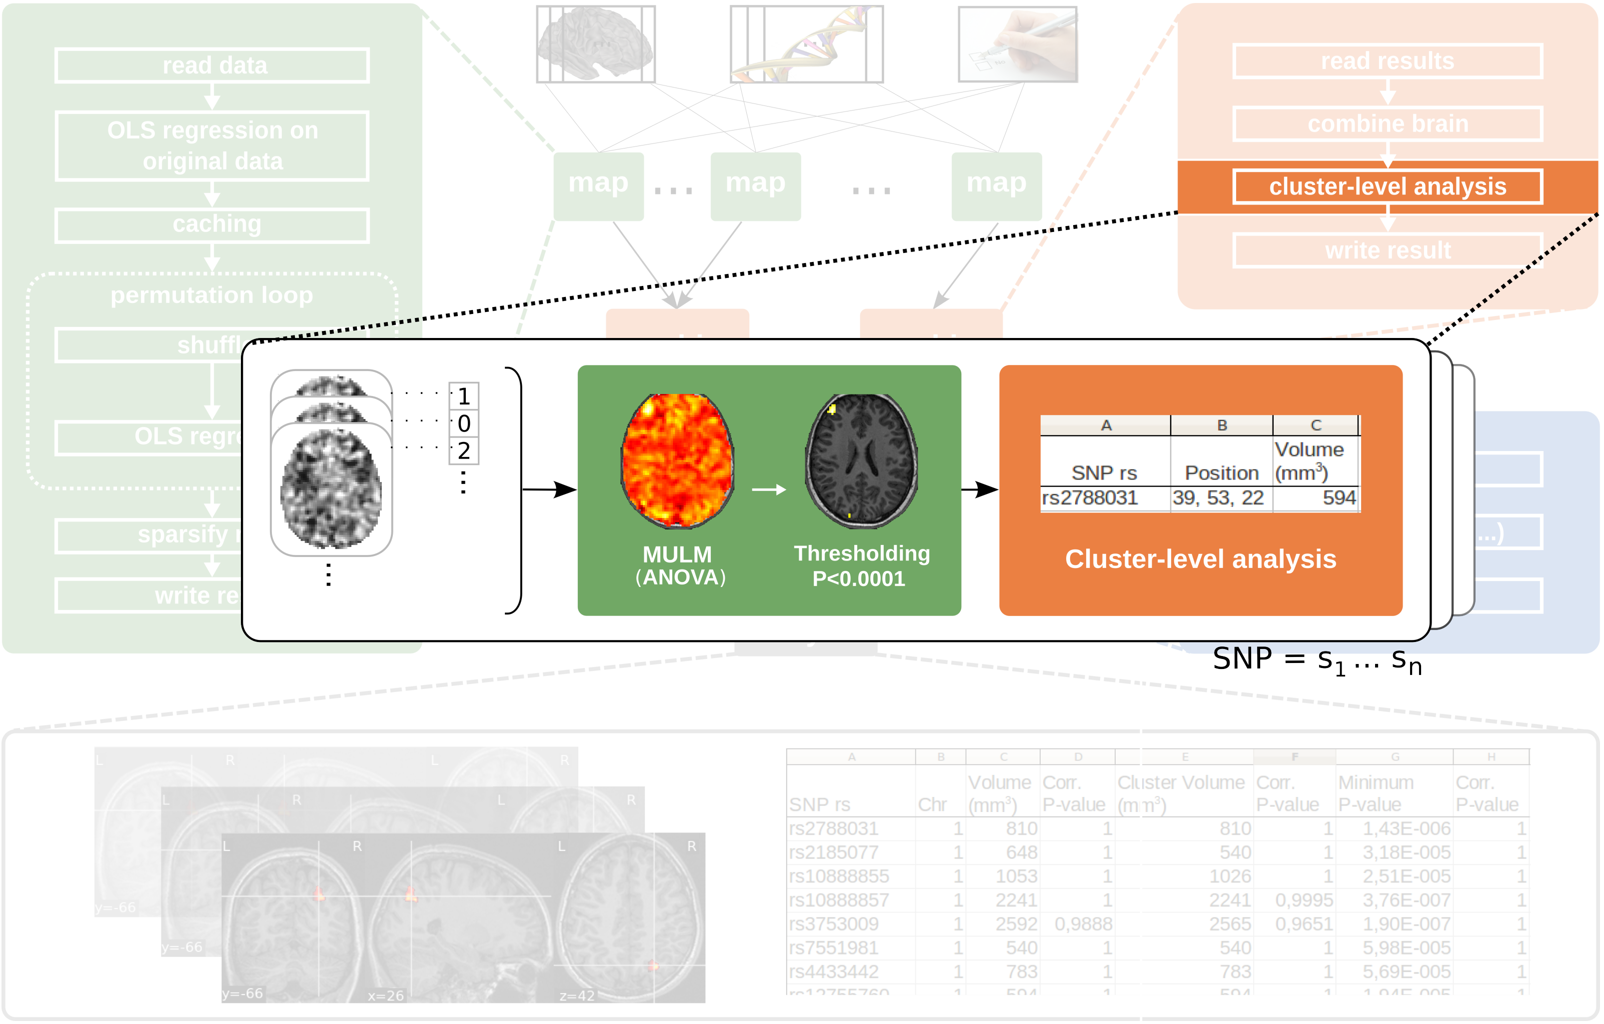

read data

OLS regression on original data

caching

permutation loop

shuffl

OLS regr

sparsify

write re

map … map … map

read results

combine brain

cluster-level analysis

write result

MULM (ANOVA)

Thresholding P<0.0001

A	B	C
SNP rs	Position	Volume (mm^3)
rs2788031	39, 53, 22	594

Cluster-level analysis

SNP = $s_1 \ldots s_n$

SNP rs	Chr	Volume (mm^3)	Corr. P-value	Cluster Volume (mm^3)	Corr. P-value	Minimum P-value	Corr. P-value
rs2788031	1	810	1	810	1	1,43E-006	1
rs2185077	1	648	1	540	1	3,18E-005	1
rs10888855	1	1053	1	1026	1	2,51E-005	1
rs10888857	1	2241	1	2241	0,9995	3,76E-007	1
rs3753009	1	2592	0,9888	2565	0,9651	1,90E-007	1
rs7551981	1	540	1	540	1	5,98E-005	1
rs4433442	1	783	1	783	1	5,69E-005	1

MAIN ACHIVEMENTS
ON THE INFRASTRUCTURE SIDE

Data-intensive Processing on Clouds: Challenges

- Computation-to-data latency is high!
- Scalable concurrent data accesses to shared data
- Need efficient Map-Reduce-like data processing
  - Hadoop is not the best we can get
  - The Reduce phase may be costly!

Scalable Storage for Processing Shared Data on Azure Clouds: TomusBlobs

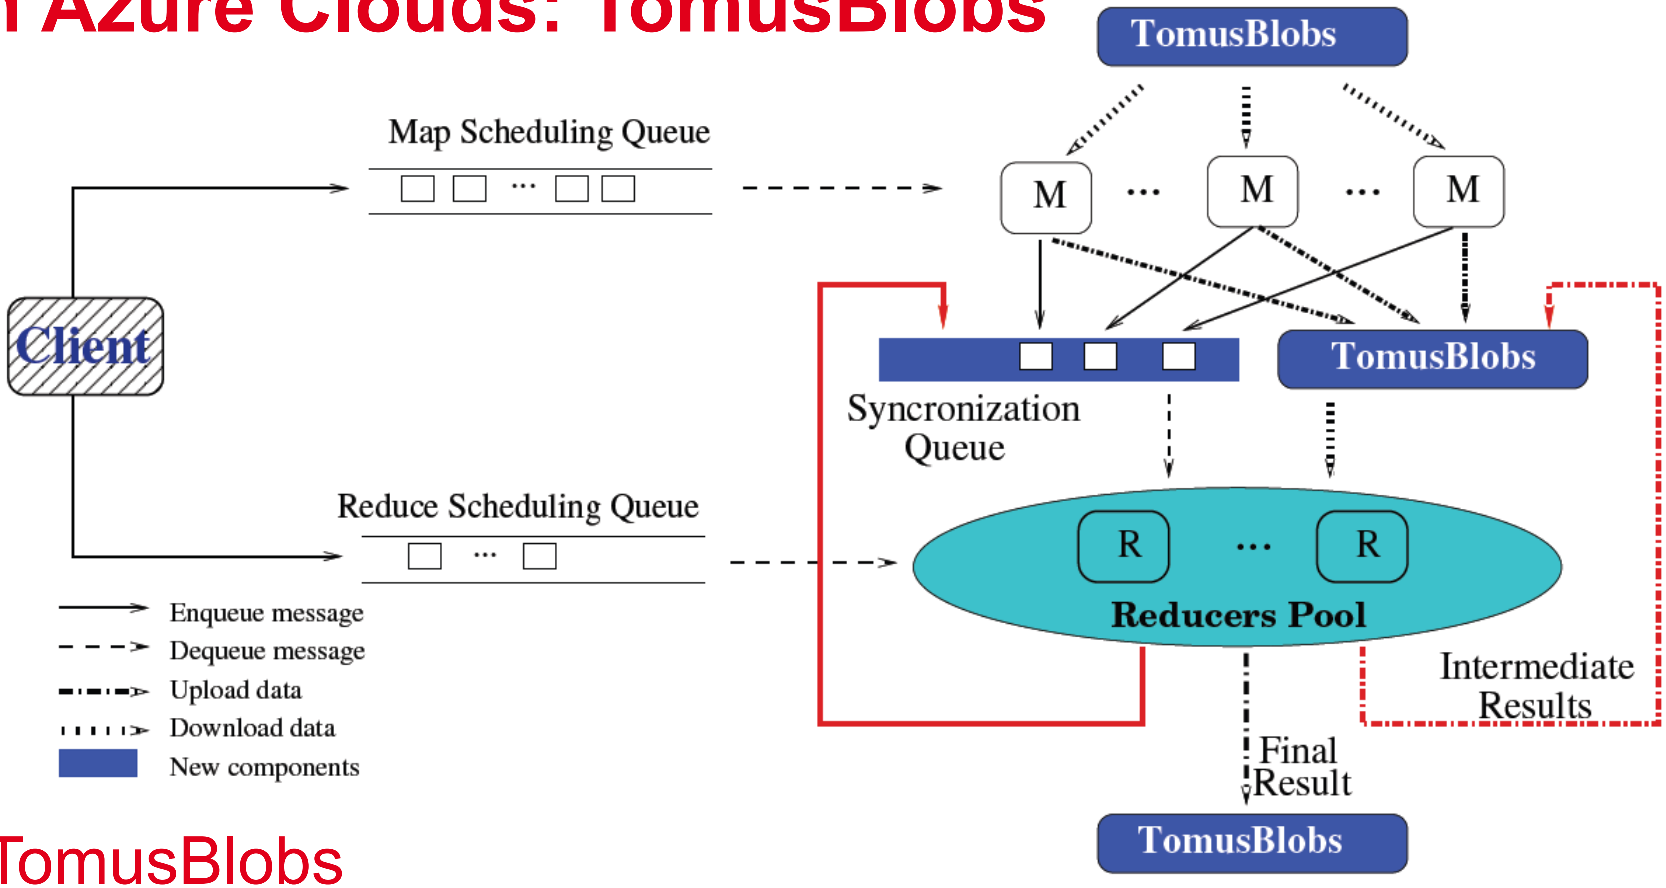

TomusBlobs

- Aggregates the virtual disks into a uniform storage
- Relies on versioning to support high throughput under heavy concurrency
  - Leverages the BlobSeer data storage software (KerData)
- Data replication

Background: BlobSeer, a Software Platform for Scalable, Distributed BLOB Management

Started in 2008, 6 PhD theses (Gilles Kahn/SPECIF PhD Thesis Award in 2011)
Main goal: optimized for concurrent accesses under heavy concurrency

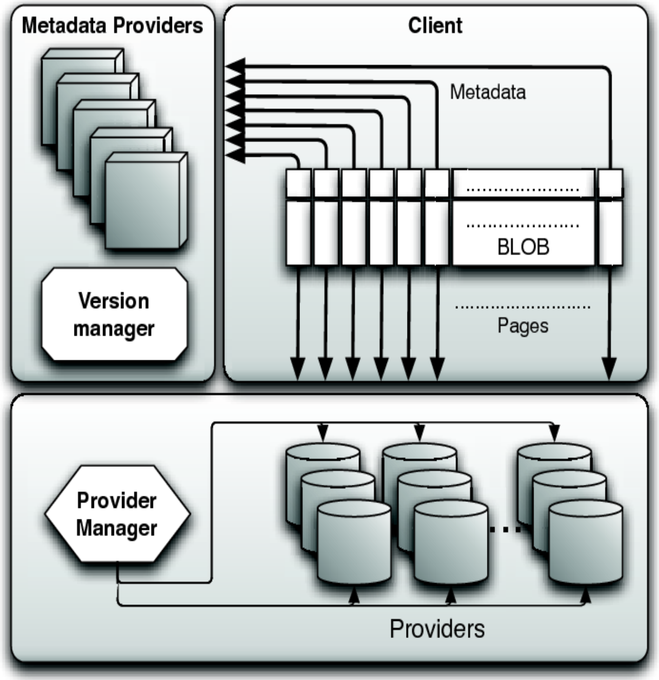

Three key ideas

Decentralized metadata management

Lock-free concurrent writes (enabled by versioning)

Write = create new version of the data

Data and metadata "patching" rather than updating

A back-end for higher-level data management systems

Short term: highly scalable distributed file systems

Middle term: storage for cloud services

Our approach

Design and implementation of distributed algorithms

Experiments on the Grid'5000 grid/cloud testbed

Validation with "real" apps on "real" platforms: Nimbus, Azure, OpenNebula clouds…

http://blobseer.gforge.inria.fr/

Using TomusBlobs for A-Brain: Results

- Gain / Azure Blobs: 45%
- Scalability: 1000 cores
- Demo available

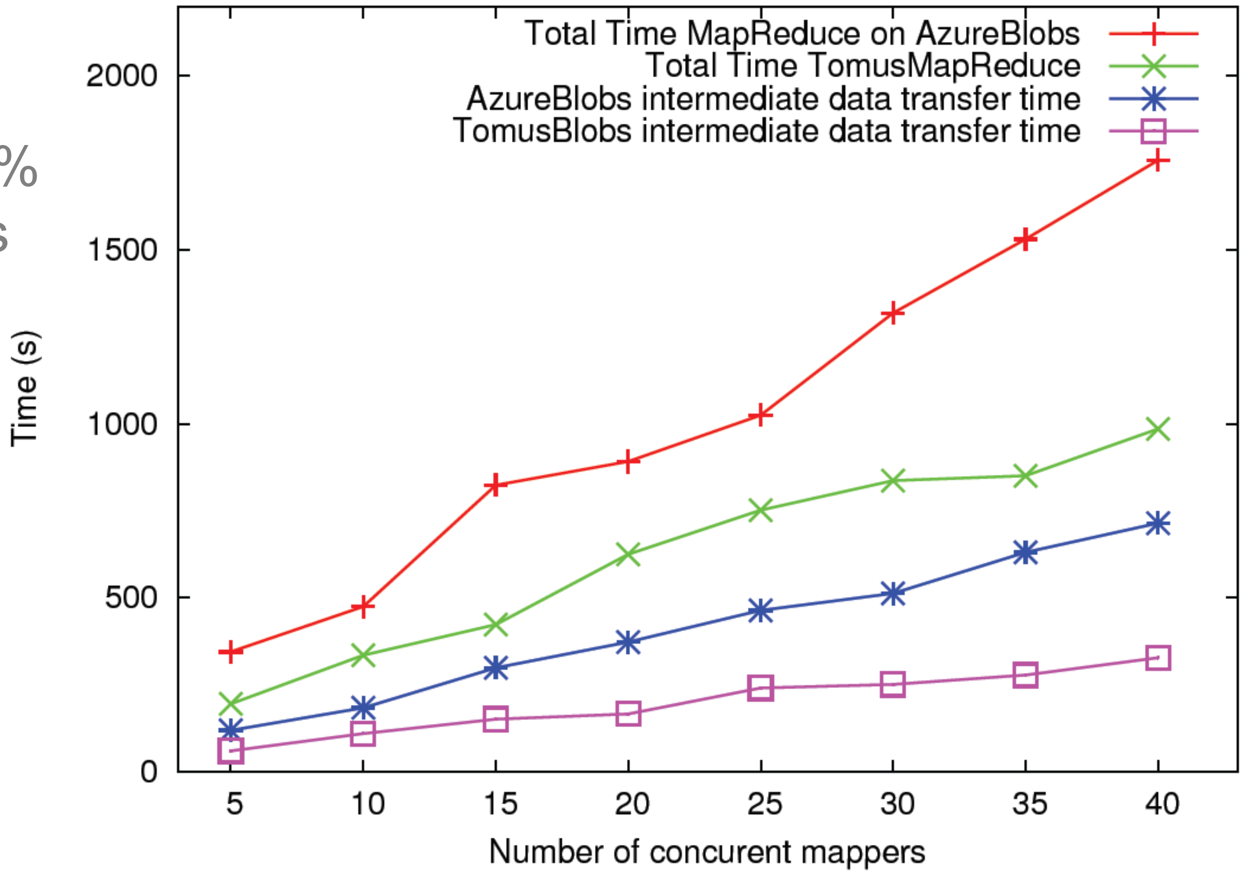

http://www.irisa.fr/kerdata/doku.php?id=abrain

Extending the MapReduce Model: MapIterativeReduce

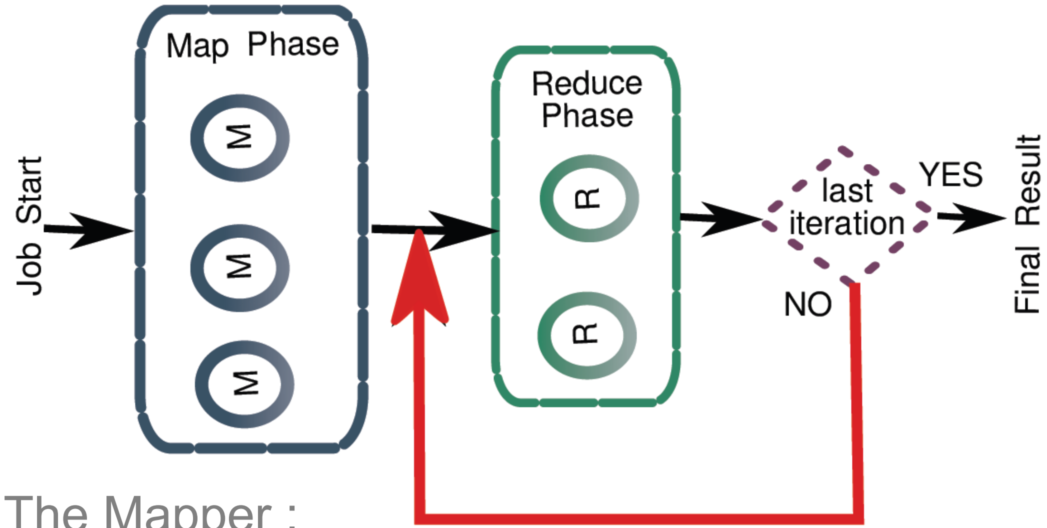

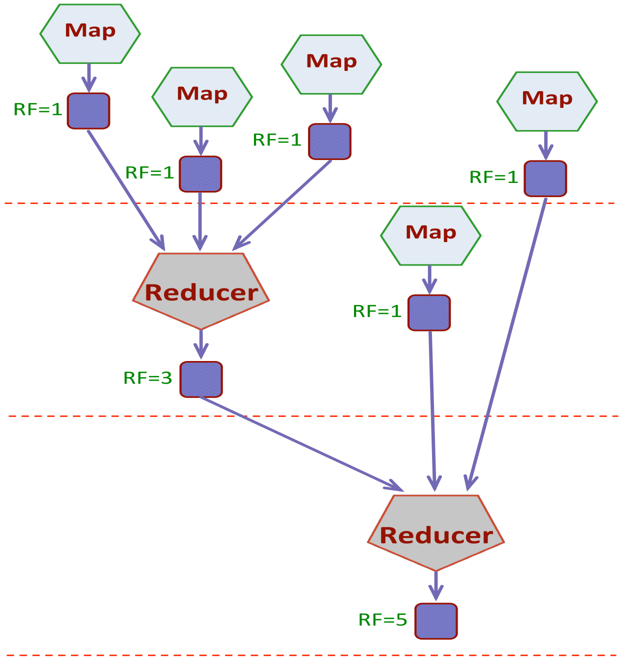

The Mapper :

- Classical map tasks

The Reducer

- Iterative reduction in two steps:
  - Receive the workload description from the Clients
  - Process intermediate results
- After each iteration, the termination condition is checked

Impact of MapIterativeReduce on A-Brain

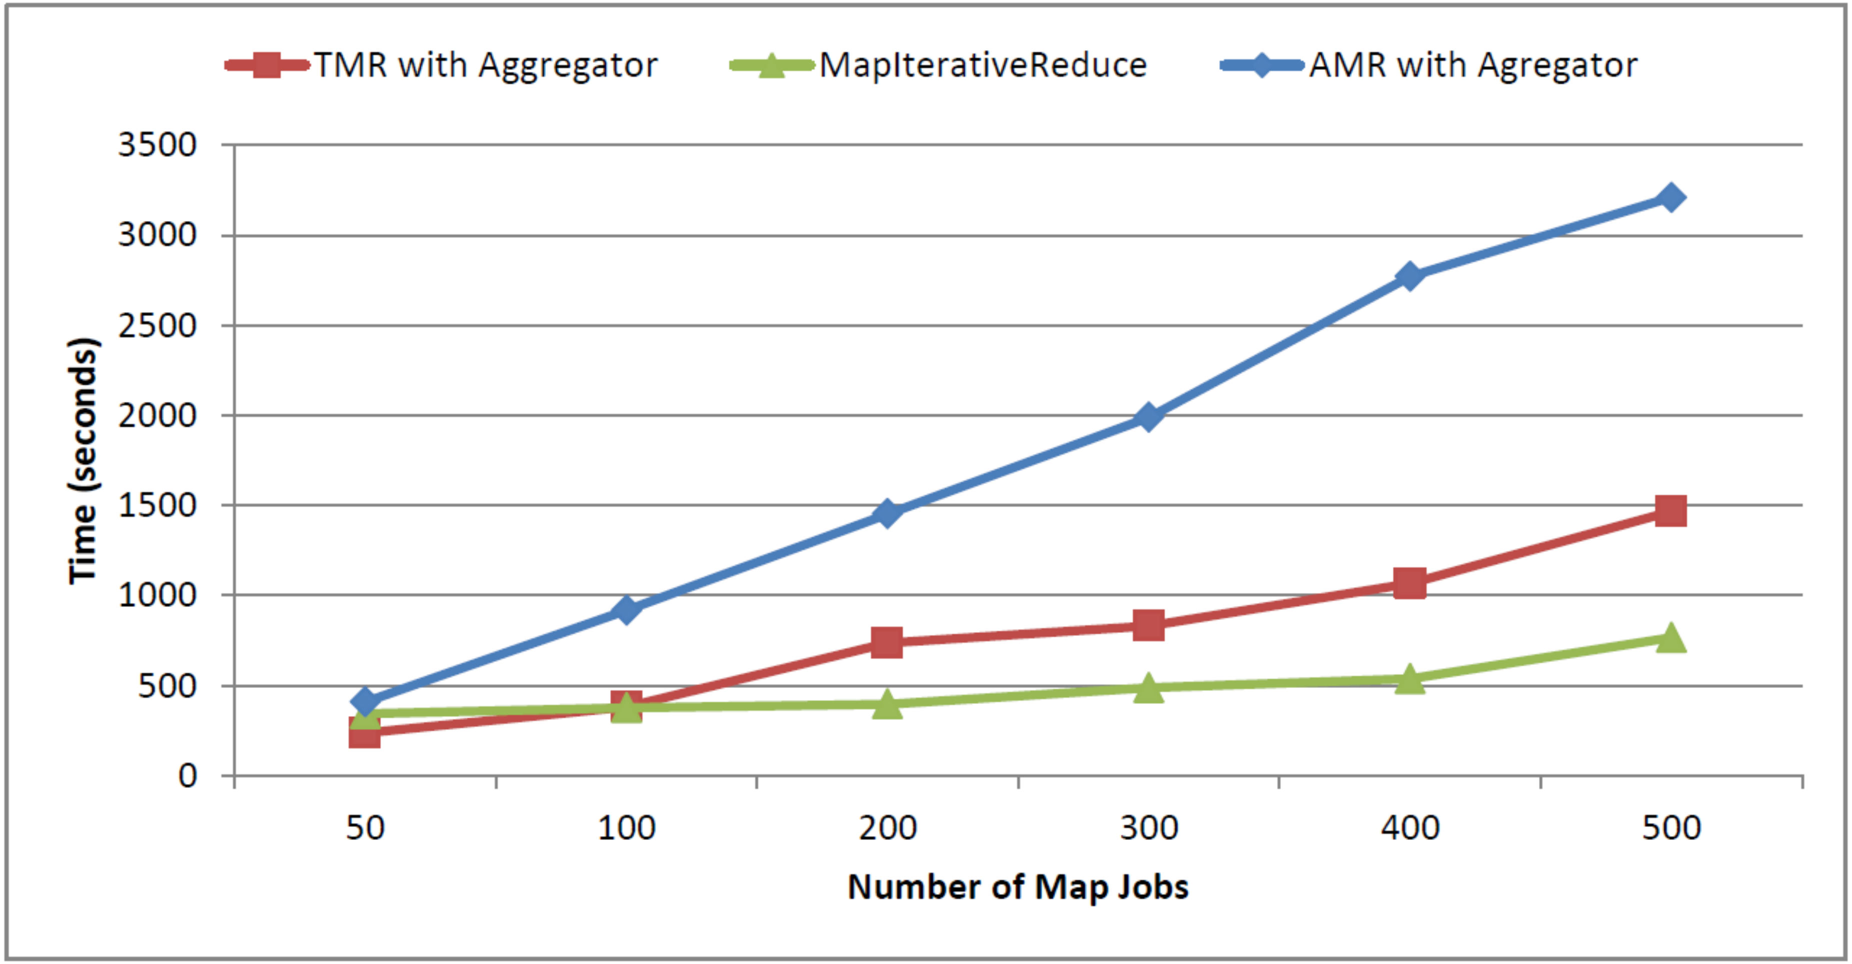

Beyond Single Site processing

- Data movements across geo-distributed deployments is costly
- Minimize the size and number of transfers
- The overall aggregate must collaborate towards reaching the goal
- The deployments work as independent services
- The architecture can be used for scenarios in which data is produced in different locations

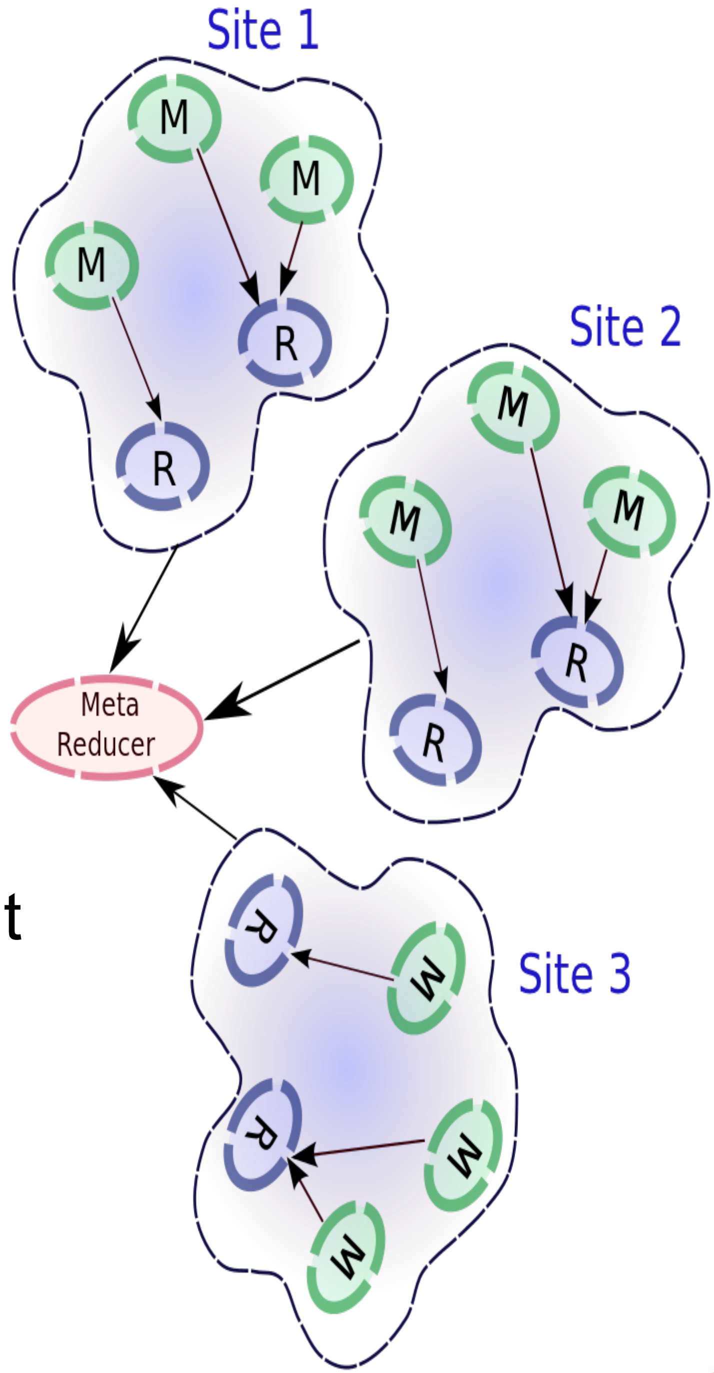

Towards Geo-distributed TomusBlobs

- TomusBlobs for intra-deployment data management
- Public Storage (Azure Blobs/ Queues) for inter-deployment communication
- Iterative Reduce technique for minimizing number of transfers (and data size)
- Balance the network bottleneck from single data center

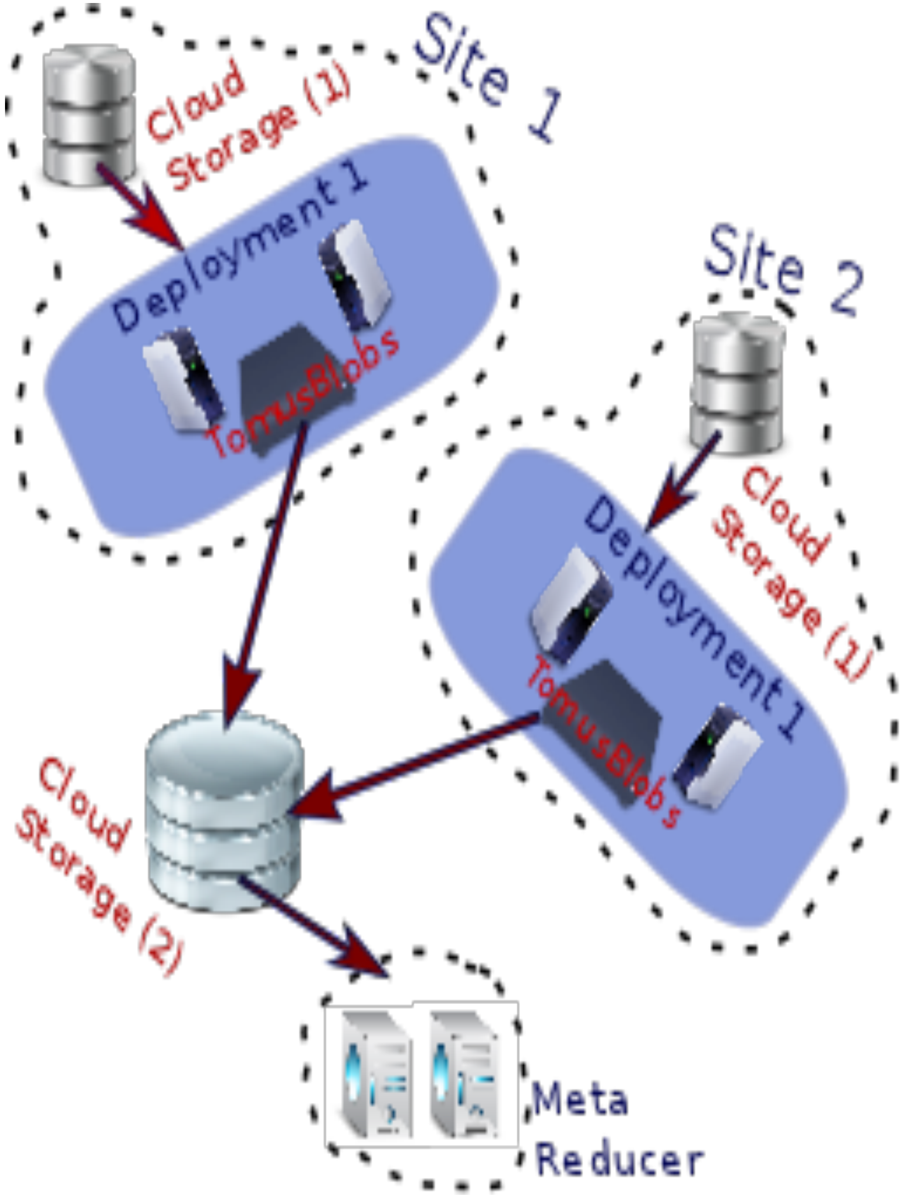

Multi-Site MapReduce

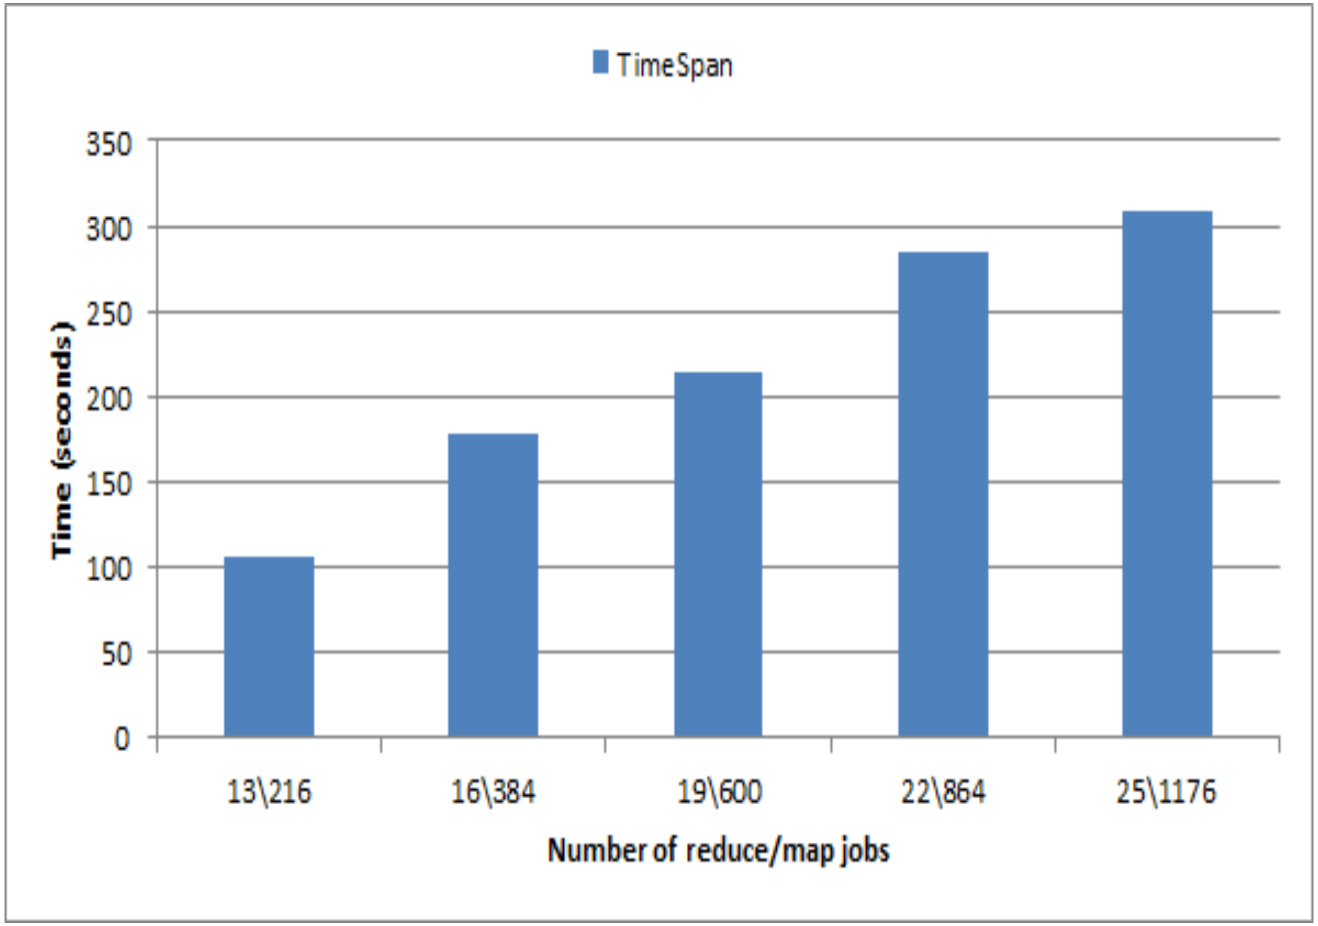

- 3 deployments (NE,WE,NUS)
- 1000 CPUs
- ABrain execution across multiple sites

MAIN ACHIVEMENTS
ON THE APPLICATION SIDE

Our contributions (0): A linear framework for mass-univariate tests

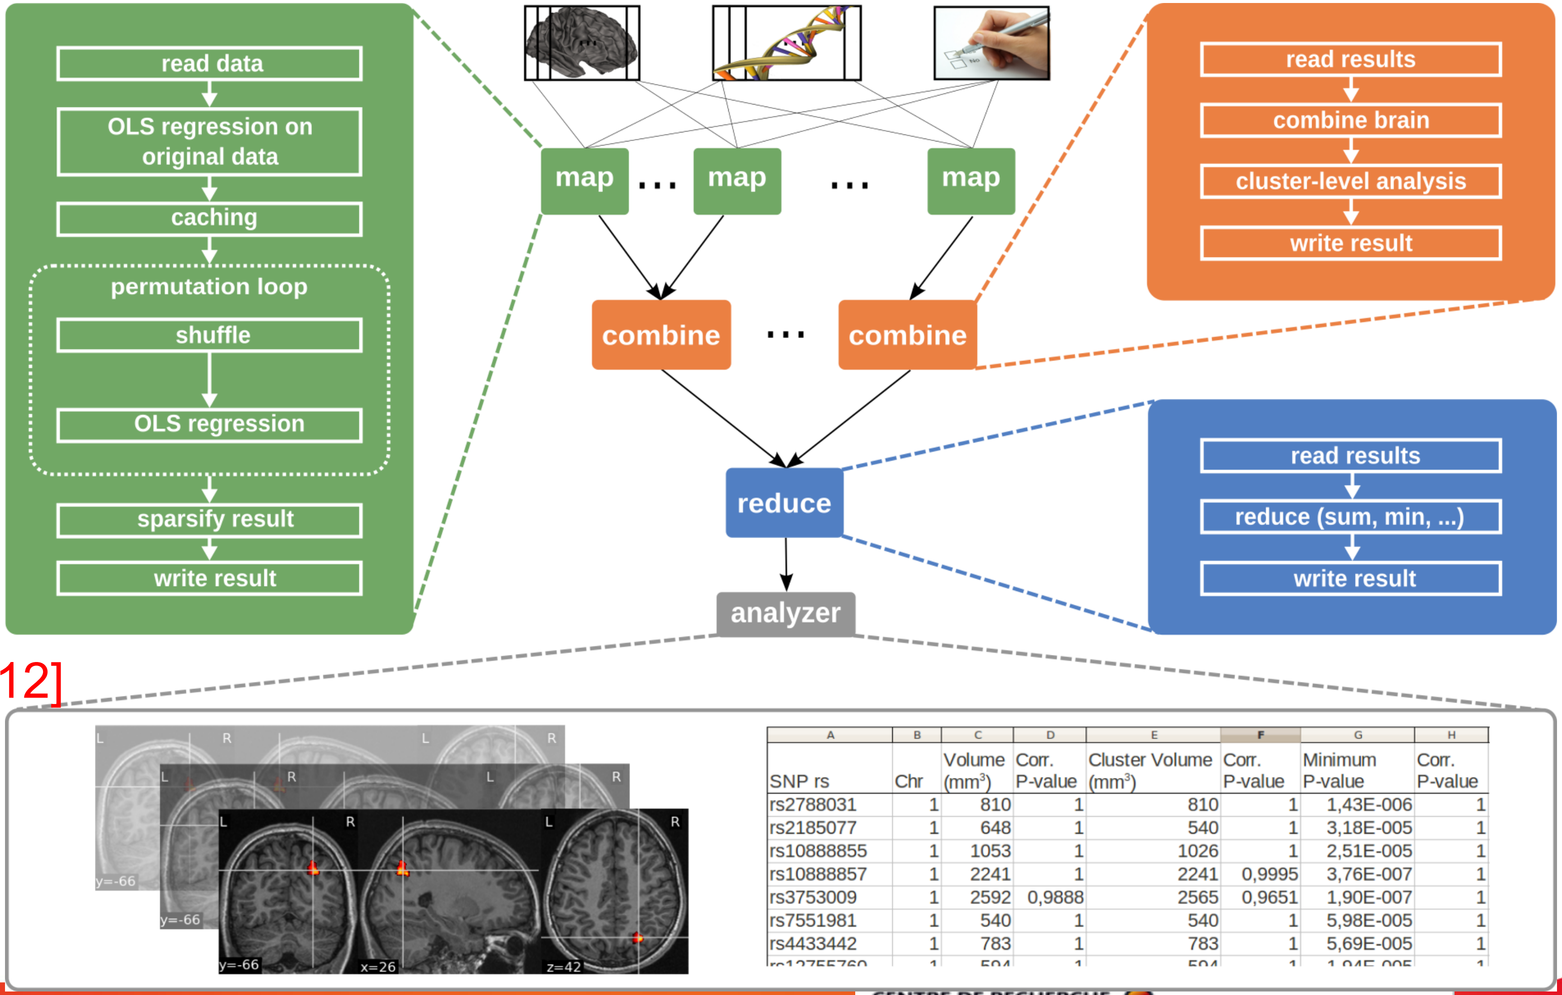

- read data
- OLS regression on original data
- caching
- permutation loop
  - shuffle
  - OLS regression
- sparsify result
- write result

map … map … map

combine … combine

- read results
- combine brain
- cluster-level analysis
- write result

reduce

- read results
- reduce (sum, min, ...)
- write result

analyzer

A	B	C	D	E	F	G	H
SNP rs	Chr	Volume (mm^3)	Corr. P-value	Cluster Volume (mm^3)	Corr. P-value	Minimum P-value	Corr. P-value
rs2788031	1	810	1	810	1	1,43E-006	1
rs2185077	1	648	1	540	1	3,18E-005	1
rs10888855	1	1053	1	1026	1	2,51E-005	1
rs10888857	1	2241	1	2241	0,9995	3,76E-007	1
rs3753009	1	2592	0,9888	2565	0,9651	1,90E-007	1
rs7551981	1	540	1	540	1	5,98E-005	1
rs4433442	1	783	1	783	1	5,69E-005	1

[Da mota et al. COMPSTAT 2012]

Our contributions (1): Improving Brain-Wide studies

- Use of a spatially regularizing prior: group features into parcels, and do the analysis on these parcels [Thirion et al. 2006]
- Remove the dependence on the parcellation choice by taking the mean across random draws

[Da Mota et al. MICCAI 2013, NeuroImage 2013]

Our contributions (1): RPBI

Randomized-parcellation based inference

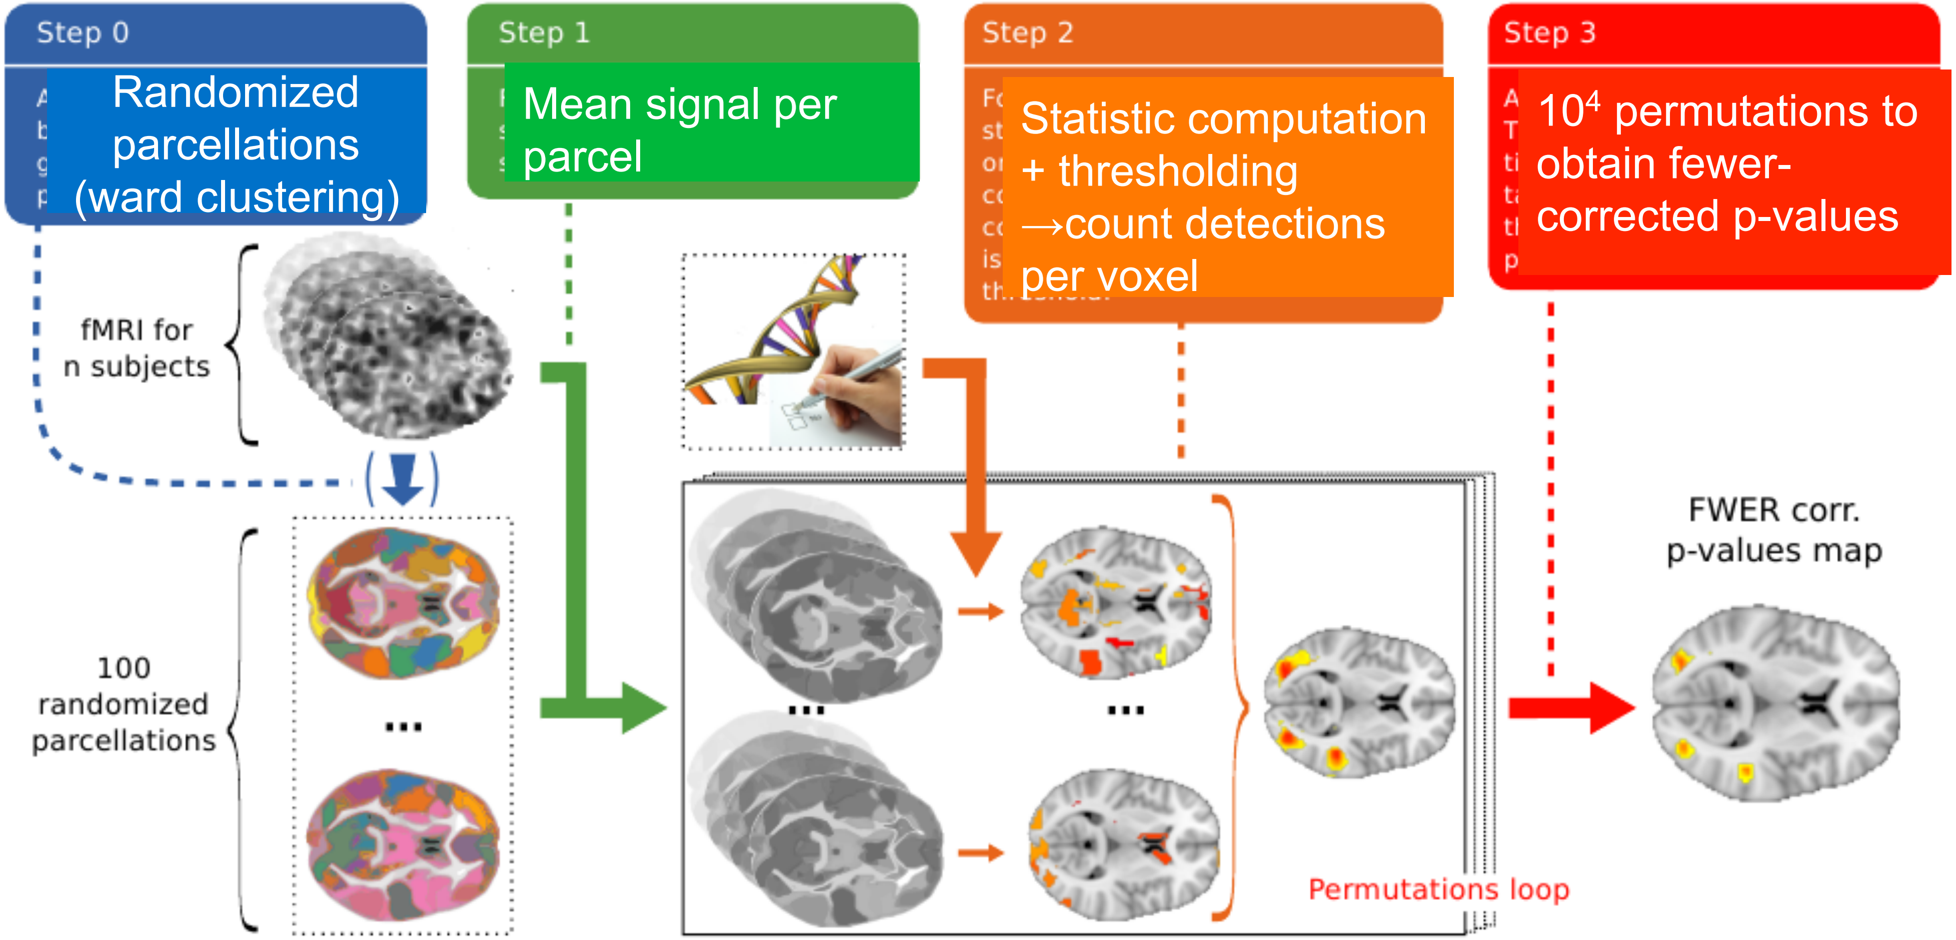

Randomized parcellations (ward clustering)

Mean signal per parcel

Statistic computation + thresholding →count detections per voxel

10^4 permutations to obtain fewer-corrected p-values

Our contributions (1): results of RPBI

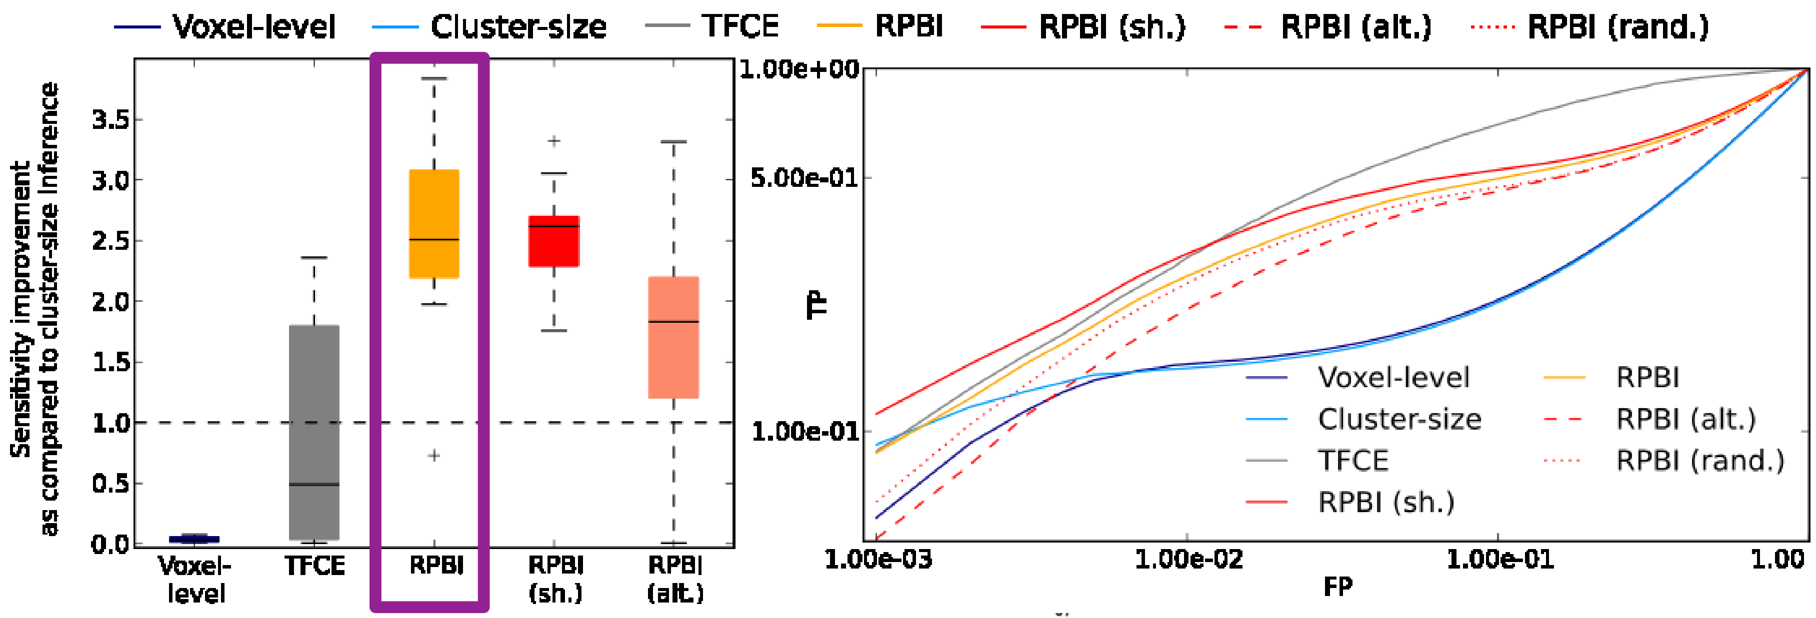

More detections on a real dataset (for a given type I error control)

More accurate model (higher ROC curves)

Higher repodusibility across groups

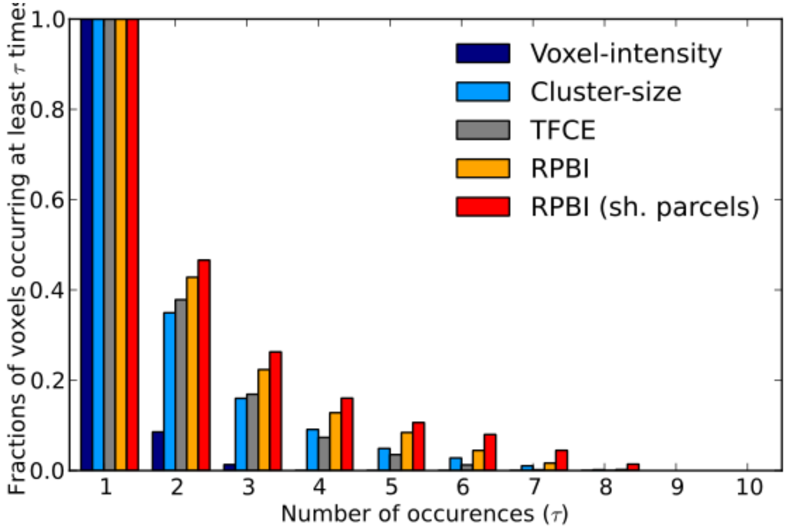

Our contributions (1): results of RPBI

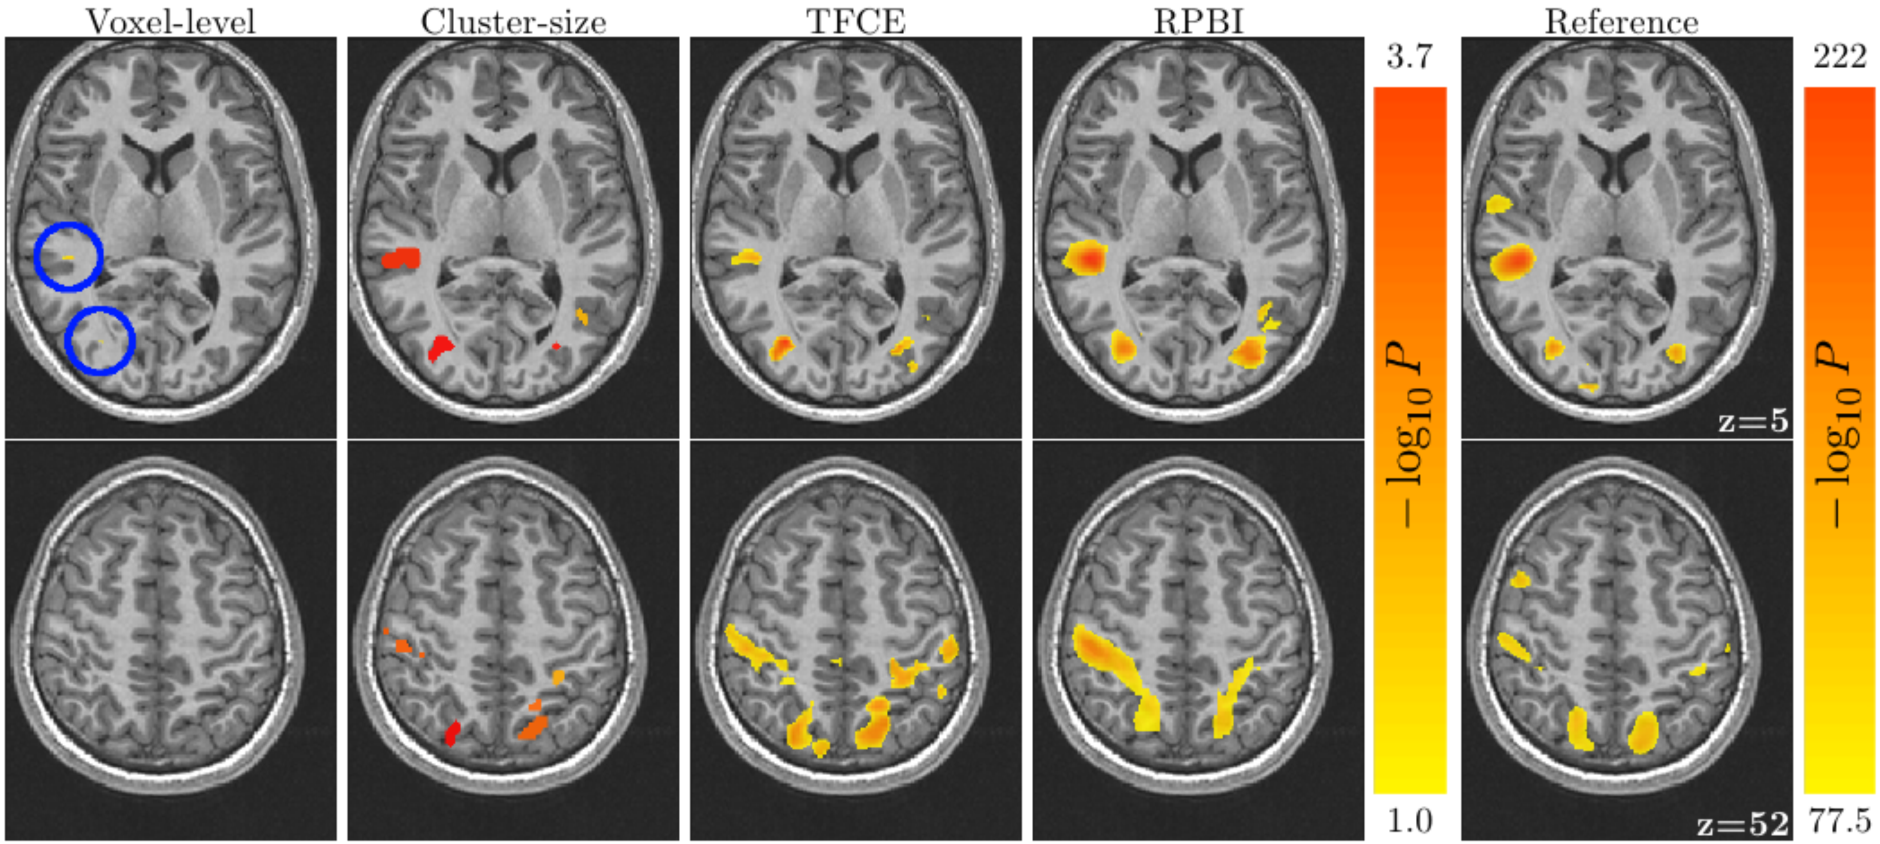

non-zero intercept test with confounds (handedness, site, sex), on an [angry faces - control] fMRI contrast from the faces protocol

Our contributions (1): results of RPBI

Experiment with a few SNPs of the ARVCF gene (close to COMT): fMRI signals upon motor response errors

RPBI uncovers a more significant association than traditional approaches

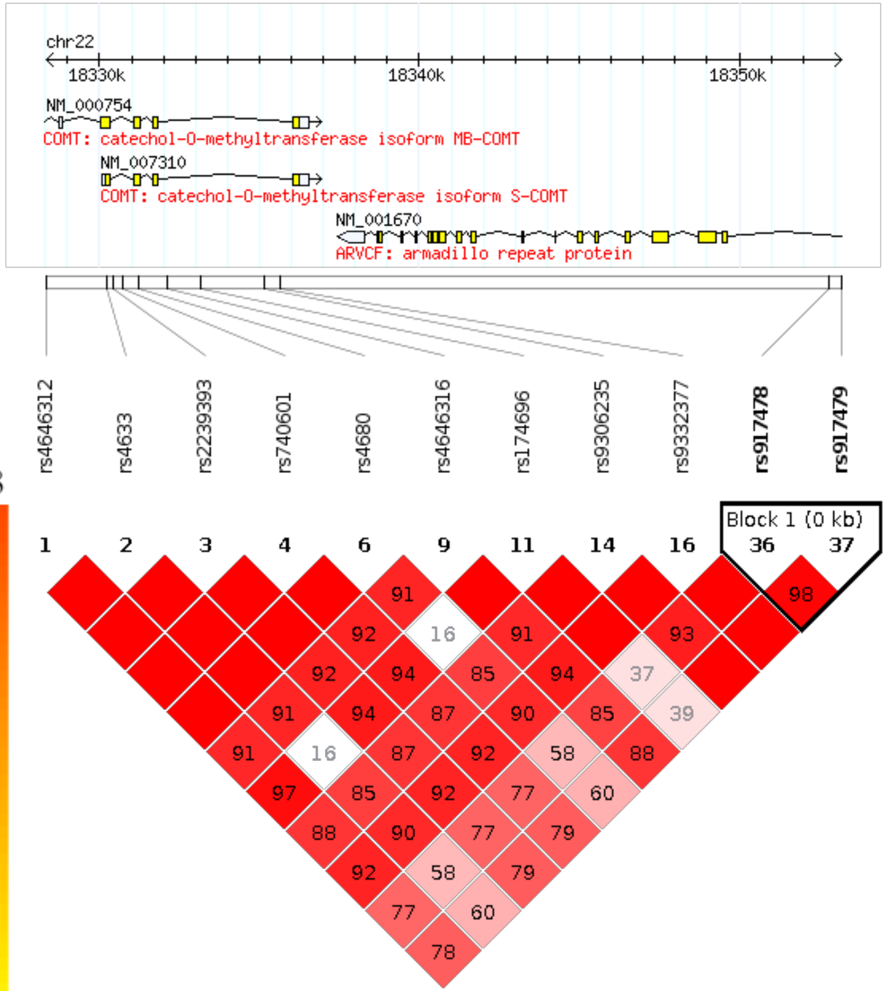

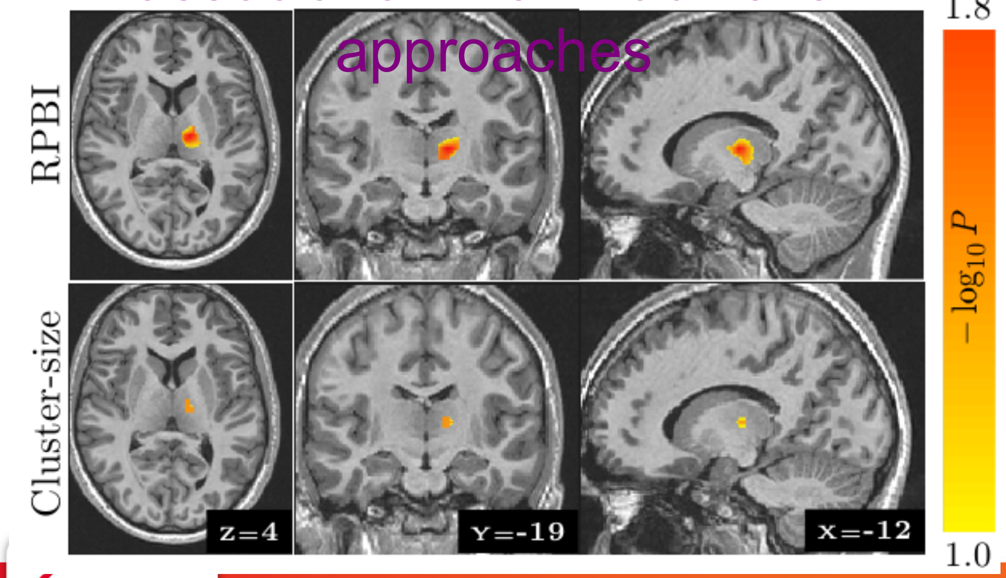

Our contributions (1): adding robustness to RPBI

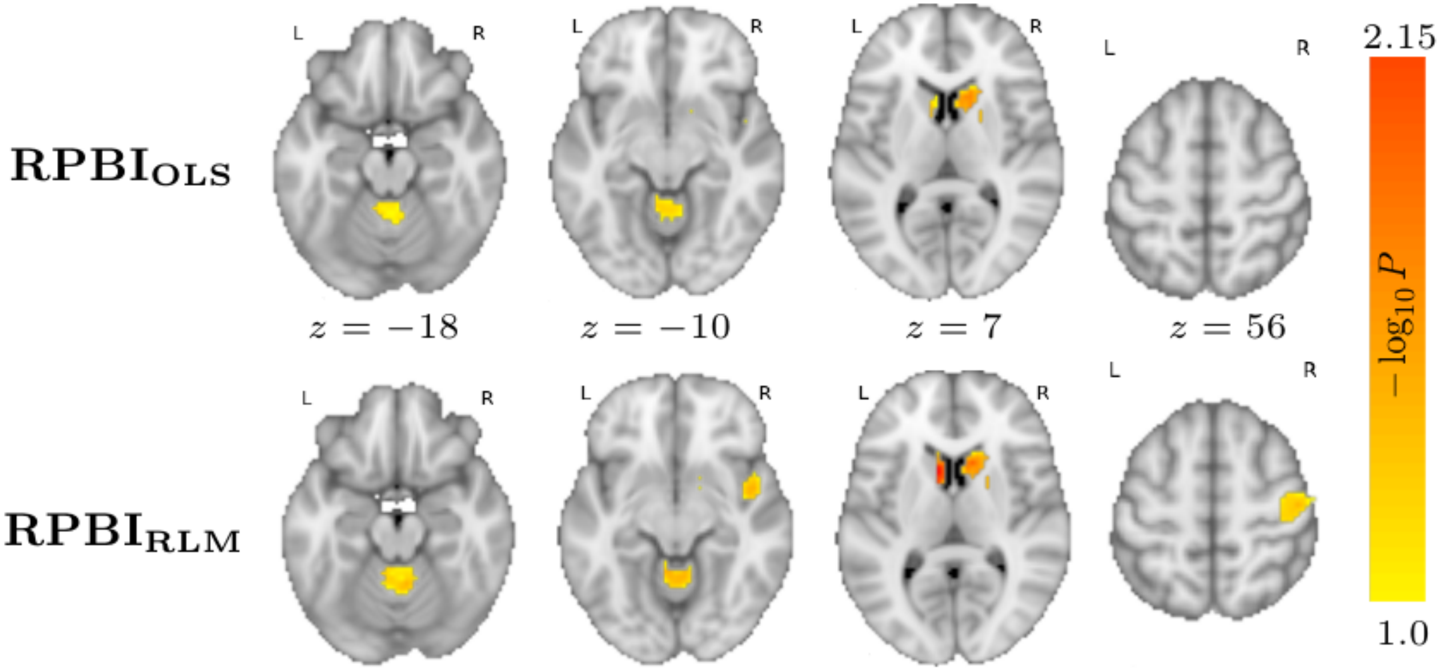

Imagen dataset: Correlation between
- the interaction of a SNP in the oxytocyn recepter gene with the number of negative life event
- the activation to angry faces

[Loth et al. 2013]

Using robust regression instead of OLS in the RPBI method yields more reliable and sometimes more sensitive detections [Fritsch et al PRNI 2013]

Our contributions (2): Improving genome-wide studies

Do not try to localize a few SNPs (among 10^6): rather assess the joint effect of all SNPs again brain variables (heritability)

- ➢ common variants are responsible of a large portion of heritability
- ➢ address the *missing variance* problem [Yang et al. Nat.gen. 2010]

Regress all the SNPs together against a given brain activation measure

FMRI signal in a subcortical region

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\beta}_1 + \boldsymbol{Z}\boldsymbol{\beta}_2 + \boldsymbol{\epsilon}$$

All SNPs

Other regressors (confounds)

[Da Mota et al. Submitted to frontiers]

Our contributions (2): Heritability estimation and test

Estimation by ridge regression

$$\hat{\beta}_1^{ridge} = \text{argmin}_{\beta_1,\beta_2 \in \mathbb{R}^p \times \mathbb{R}^q} \left\{ \|\mathbf{Y} - \mathbf{X}\beta_1 - \mathbf{Z}\beta_2\|^2 + \lambda \|\beta_1\|^2 \right\}$$

λ is learned by cross-validation

Test = amount of explained variance in a cross-validation scheme

$$CV\text{-}R^2 = 1 - mean_{(train,test)\in \text{split(n)}} \frac{\|\boldsymbol{Y}^{test} - \boldsymbol{X}^{test}\boldsymbol{\beta}_1^{train} - \boldsymbol{Z}^{test}\boldsymbol{\beta}_2^{train}\|^2}{\|\boldsymbol{Y}^{test} - \boldsymbol{Z}^{test}\boldsymbol{\beta}_2^{train}\|^2}$$

Average Predictive explained variance = a proxy for heritability

Our contributions (2):
Results with heritability

ROI name		$CV\text{-}R^2$	fwe corr. p-value
Thalamus	left	**0.026**	1.10^{-4}
	right	**0.038**	1.10^{-4}
Caudate	left	**0.003**	2.10^{-4}
	right	-0.012	3.10^{-4}
Putamen	left	**0.019**	1.10^{-4}
	right	**0.006**	2.10^{-4}
Pallidum	left	**0.018**	1.10^{-4}
	right	-0.010	3.10^{-4}
Hippocampus	left	**0.010**	2.10^{-4}
	right	**0.020**	1.10^{-4}
Amygdala	left	**0.016**	1.10^{-4}
	right	**0.015**	1.10^{-4}
Accumbens	left	**0.022**	1.10^{-4}
	right	-0.002	2.10^{-4}

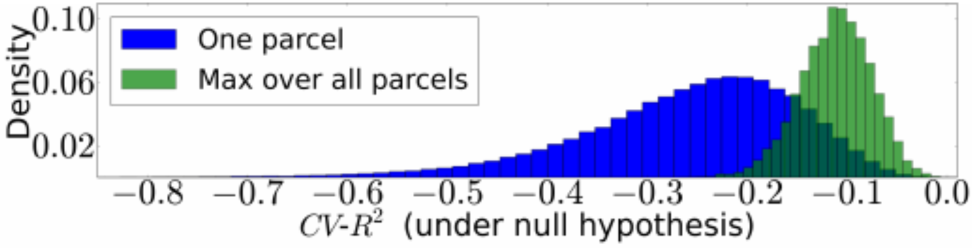

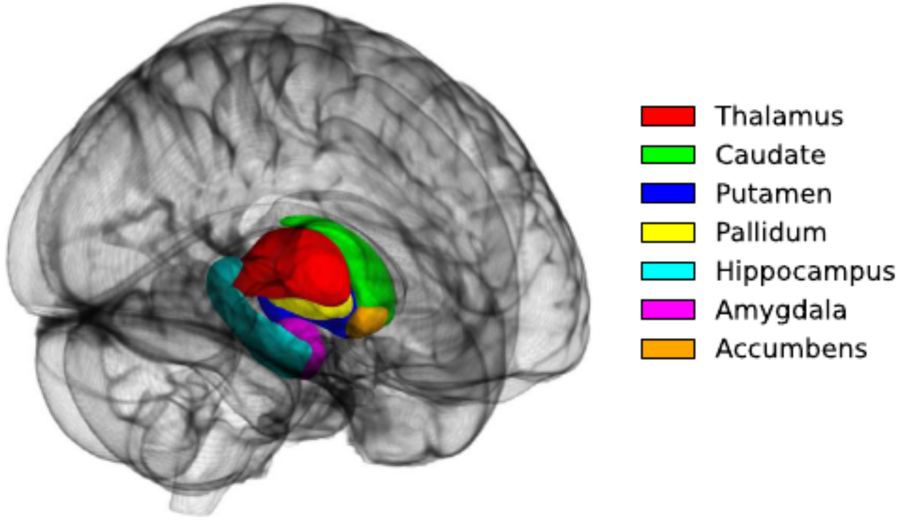

Experiment on the Imagen dataset: heritability of the stop failure brain activation signals in the sub-cortical nuclei:The signals are significantly more heritable than chance in all regions considered

Conclusion: where we are

- Good method for brain-wide association RPBI
- Genome-wide associations: build on the ridge-based heritability estimate
  - Analysis at the level of pathways, genes
  - Robust version of ridge regression ?
- Application:
  - Not enough data !
  - need more precise hypotheses to test
  - Need more feature engineering

Conclusion: what we learned from A-brain

- Using the cloud can be advantageous:
  - Do not need to own the cluster
  - Resources owned until the end of the computation
  - Ease of use: execute the same code as the usual one
- Progress still needed to get closer to the power of a bare cluster

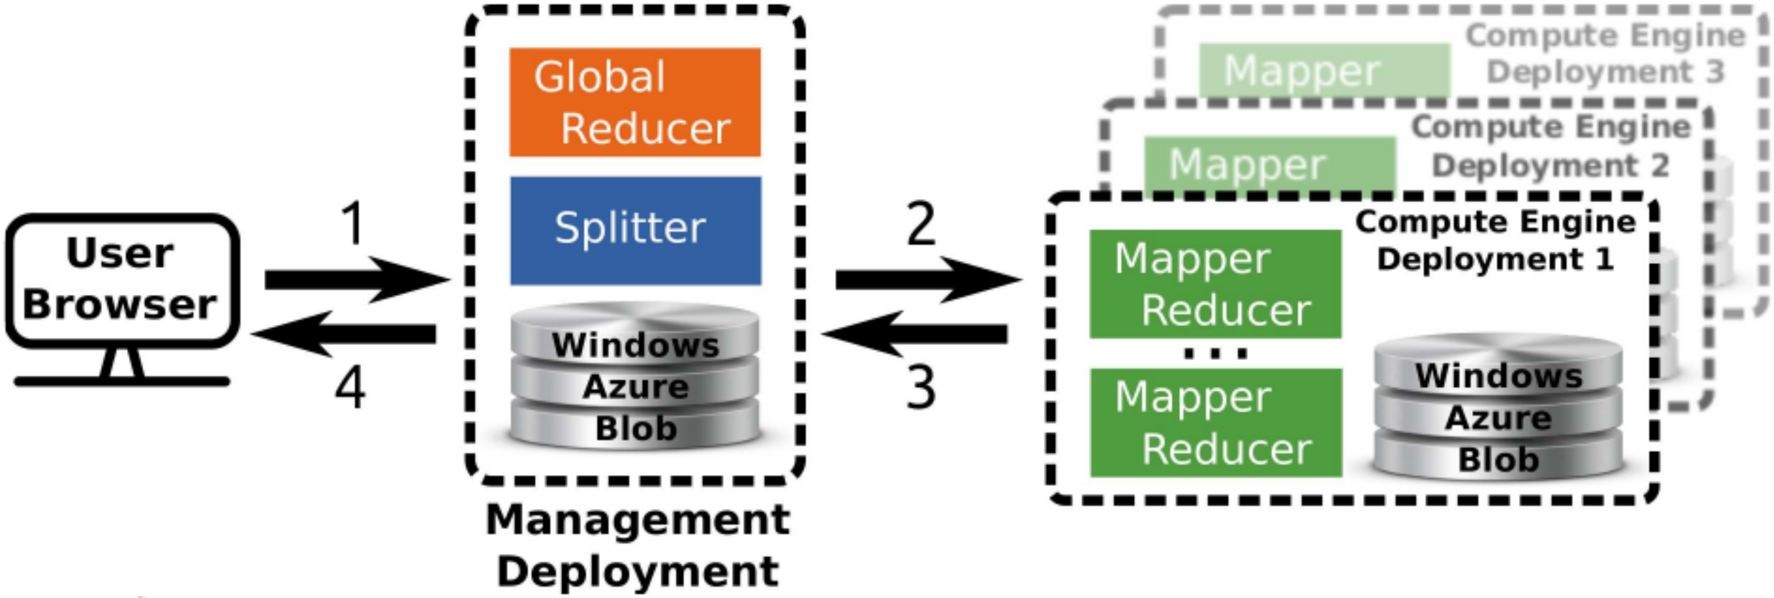

Two Things to Take Away

- The TomusBlobs data-storage layer developed within the A-Brain project was demonstrated to scale up to 1000 cores on 3 Azure data centers.
  - It exhibits improvements in execution time up to 50% compared to standard solutions based on Azure BLOB storage.

- The consortium has provided the first statistical evidence of the heritability of functional signals in a failed stop task in basal ganglia, using a ridge regression approach, while relying on the Azure cloud to address the computational burden.

Publications

Journals

- Alexandru Costan, Radu Tudoran, Gabriel Antoniu, Goetz Brasche. TomusBlobs : Scalable Data-intensive Processing on Azure Clouds. Concurrency and Computation Practice and Experience, Wiley, 2013. URL: http://onlinelibrary.wiley.com/doi/10.1002/cpe.3034/abstract.

- Benoit Da Mota, Virgile Fritscha, Gaël Varoquaux, Tobias Banaschewski, Gareth J. Barker , Arun L.W. Bokde, Uli Bromberg , Patricia Conrod, Jürgen Gallinat, Hugh Garavan, Jean-Luc Martinot, Frauke Nees, Tomas Pausl, Zdenka Pausova , Marcella Rietschel, Michael N. Smolka, Andreas Ströhle, Vincent Frouin, Jean-Baptiste Poline, Bertrand Thirion, the IMAGEN consortium. Randomized Parcellation Based Inference. NeuroImage, Elsevier, in Press.

- Benoit Da Mota, Radu Tudoran, Alexandru Costan, Gael Varoquaux, Goetz Brasche, Patricia Conrod, Herve Lemaitre, Tomas Paus, Marcella Rietschel, Vincent Frouin, Jean-Baptiste Poline, Gabriel Antoniu, Bertrand Thirion and the IMAGEN Consortium. Machine Learning Patterns for Neuroimaging-Genetic Studies in the Cloud. Submitted to Frontiers in the Cloud.

Electronic Journals

- Gabriel Antoniu, Alexandru Costan, Benoit Da Mota, Bertrand Thirion, Radu Tudoran. A-Brain: Using the Cloud to Understand the Impact of Genetic Variability on the Brain. **ERCIM News**, April 2012.

Publications

Conferences and workshops (2013)

- Radu Tudoran, Alexandru Costan, Ramin Rezai Rad, Goetz Brasche and Gabriel Antoniu. Adaptive File Management for Scientific Workflows on the Azure Cloud. IEEE International Conference on Big Data (IEEE BigData 2013), October 6-9, 2013, Santa Clara, CA, USA. **Acceptance rate: 17%**.

- Radu Tudoran, Alexandru Costan, Gabriel Antoniu. DataSteward : Using Dedicated Compute Nodes for Scalable Data Management on Public Clouds. In Proc. of ISPA 2013- 11th IEEE International Symposium on Parallel and Distributed Processing with Applications, Melbourne, Australia, July 2013.

- Benoit da Mota, Virgile Fritsch, Gaël Varoquaux, Vincent Frouin, Jean-Baptiste Poline, and Bertrand Thirion. Distributed High-Dimensional Regression with Shared Memory for Neuroimaging-Genetic Studies. in Euroscipy 2013.

- Benoit Da Mota, Virgile Fritsch, Gaël Varoquaux, Vincent Frouin, Jean-Baptiste Poline, and Bertrand Thirion. Enhancing the Reproducibility of Group Analysis with Randomized Brain Parcellations. In MICCAI - 16th International Conference on Medical Image Computing and Computer Assisted Intervention - 2013, Nagoya, Japan, June 2013.

- Virgile Fritsch, Benoit Da Mota, Gaël Varoquaux, Vincent Frouin, Eva Loth, Jean-Baptiste Poline and Bertrand Thirion. Robust Group-Level Inference in Neuroimaging Genetic Studies. In Pattern Recognition in Neuroimaging, Philadelphie, United States, May 2013.

Publications

Conferences and workshops (2012)

- Radu Tudoran, Alexandru Costan, Gabriel Antoniu, Hakan Soncu. “TomusBlobs: Towards Communication-Efficient Storage for MapReduce Applications in Azure.” In Proc. 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (**CCGrid'2012**), May 2012, Ottawa, Canada.

- Radu Tudoran, Alexandru Costan, Gabriel Antoniu, Luc Bougé. A Performance Evaluation of Azure and Nimbus Clouds for Scientific Applications. In Proc. **CloudCP** 2012 - 2nd International Workshop on Cloud Computing Platforms, Held in conjunction with the ACM SIGOPS Eurosys 12 conference, Apr 2012, Bern, Switzerland.

- Radu Tudoran, Alexandru Costan, Benoit Da Mota, Gabriel Antoniu, Bertrand Thirion. A-Brain: Using the Cloud to Understand the Impact of Genetic Variability on the Brain. 2012 **Cloud Futures** Workshop, Berkeley, May 2012.

- Radu Tudoran, Alexandru Costan, Gabriel Antoniu. MapIterativeReduce: A Framework for Reduction-Intensive Data Processing on Azure Clouds. Third International Workshop on MapReduce and its Applications (**MAPREDUCE'12**), held in conjunction with ACM HPDC'12., Jun 2012, Delft, Netherlands.

- Benoit Da Mota, Vincent Frouin, Edouard Duchesnay, Soizic Laguitton, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. A fast computational framework for genome-wide association studies with neuroimaging data. 20th International Conference on Computational Statistics (**COMPSTAT 2012**), Aug 2012, Lamissol, Cyprus.

- Benoit Da Mota, Michael Eickenberg, Soizic Laguittton, Vincent Frouin, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. A MapReduce Approach for Ridge Regression in Neuroimaging-Genetic Studies. Data- and Compute-Intensive Clinical and Translational Imaging Applications Workshop **(DCICTIA-MICCAI'12**), held in conjunction with the 15th International Conference on Medical Image Computing and Computer Assisted Intervention, Oct 2012, Nice, France.