

The Impact Of Hardware Variability On Applications Packaged With Docker And Guix: A Case Study In Neuroimaging

Gaël Vila ¹
 Emmanuel Medernach ²
 Inés Gonzalez Pepe ³
 Axel Bonnet ¹
 Yohan Chatelain ³
 Michaël Sdika ¹
 Tristan Glatard ³
 Sorina Camarasu-Pop ¹

¹CREATIS, CNRS (UMR 5220), INSERM (U1294), INSA Lyon, Université de Lyon, France

²IPHC, CNRS/IN2P3, Université de Strasbourg, Strasbourg, France

³Department of Computer Science and Software Engineering, Concordia University Montreal, Canada



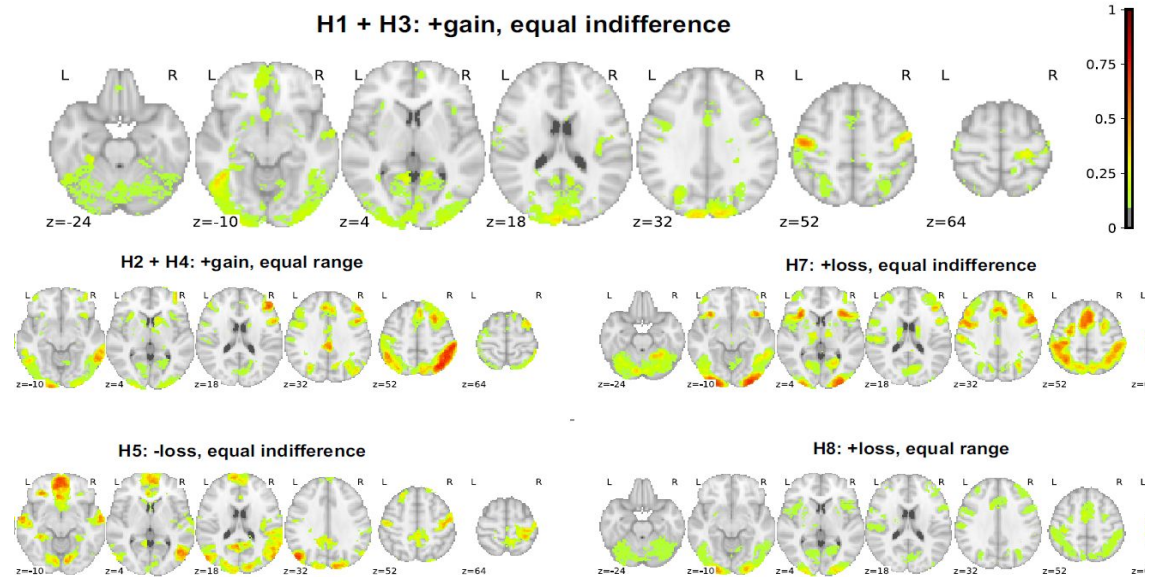
Reproducibility Issues in Neuroimaging

Setup

- **1** dataset
- **70** teams
- **9** hypotheses




Findings

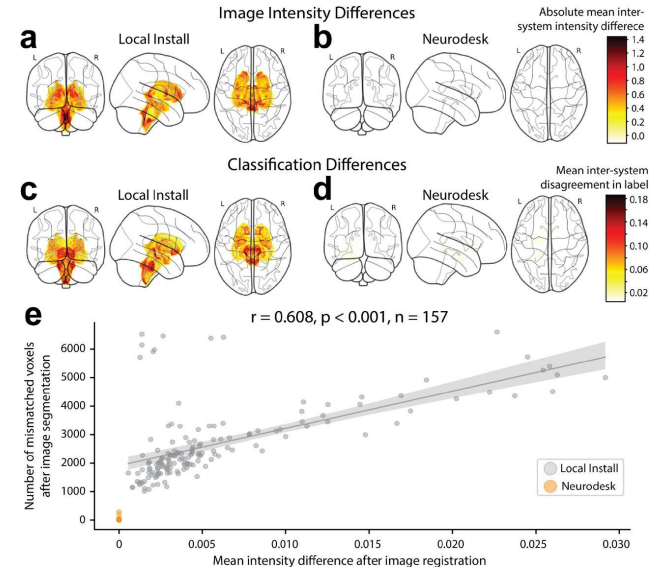
- Analytical flexibility
- Variability of results
- Optimism bias



[R. Botvinik-Nezer et al](#), "Variability in the analysis of a single neuroimaging dataset by many teams" *Nature* 2020

Computational reproducibility

- Main causes
 - Software dependencies and their evolution over time
 - Numerical instability due to floating point arithmetic
- Containerization  
 - Package and run an application and its dependencies
- Guix 
 - Functional package manager
 - Reproducible computational environments

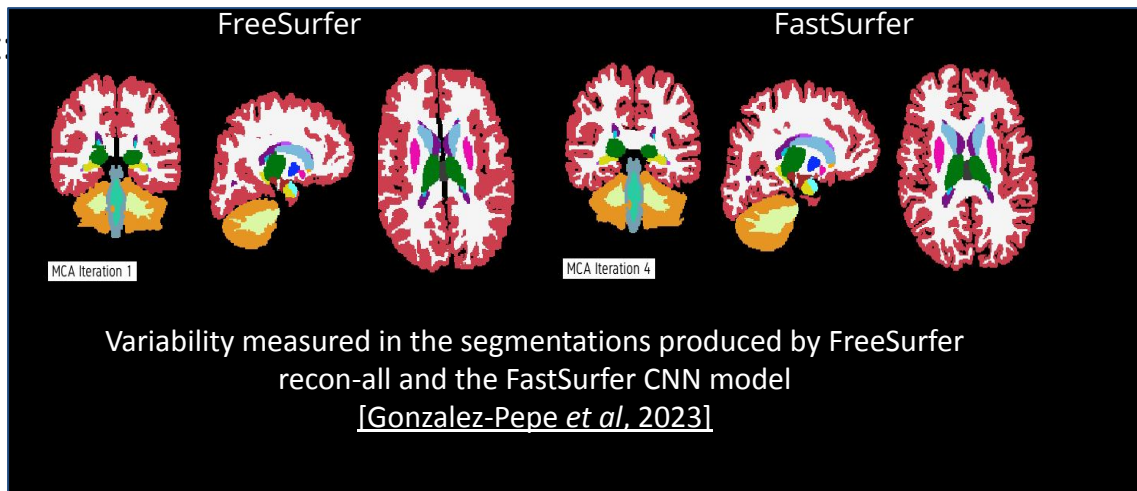



What about **hardware heterogeneity**?

[A.I. Renton et al](#), *Neurodesk*
Nature 2024

Numerical instability and MCA

- Floating point arithmetic
 - Approximate real numbers within a limited precision => rounding errors
- Monte Carlo Arithmetic (MCA)
 - Noise injection into floating-point operations: $inexact(x) = x + 2^{e_x - t} \xi$
- Random Rounding (RR)
 - Perturbs the function output
- MCA Implementations
 - [Verificarlo](#)
 - [Fuzzy Github repository](#)




Objectives


- Evaluate the **impact of hardware variability**
- Compare and correlate hardware variability to
 - **Software variability** encountered in different software packages
 - **Numerical variability** resulting from MCA RR



original image



reference



registered image

FSL FLIRT

- FMRIB Software Library ([FSL](#))
 - Library of analysis tools for FMRI, MRI and diffusion brain imaging data
- FLIRT: FMRIB's Linear Image Registration Tool
 - Affine brain registration: align a brain scan with another one through rotation, translation, scaling and shearing
- FLIRT outputs
 - Registered brain image in NIfTI format (.nii.gz)
 - Transformation matrix in text format (.mat)

```
1.129633431    0.009161432163   -0.002279976965   -2.097511242
-0.004720817456  1.028899087     0.3437343964     -50.46994368
0.01111236612  -0.413128704     1.142416095      -28.62331337
0              0              0              1
```

Example of transformation matrix (.mat file)

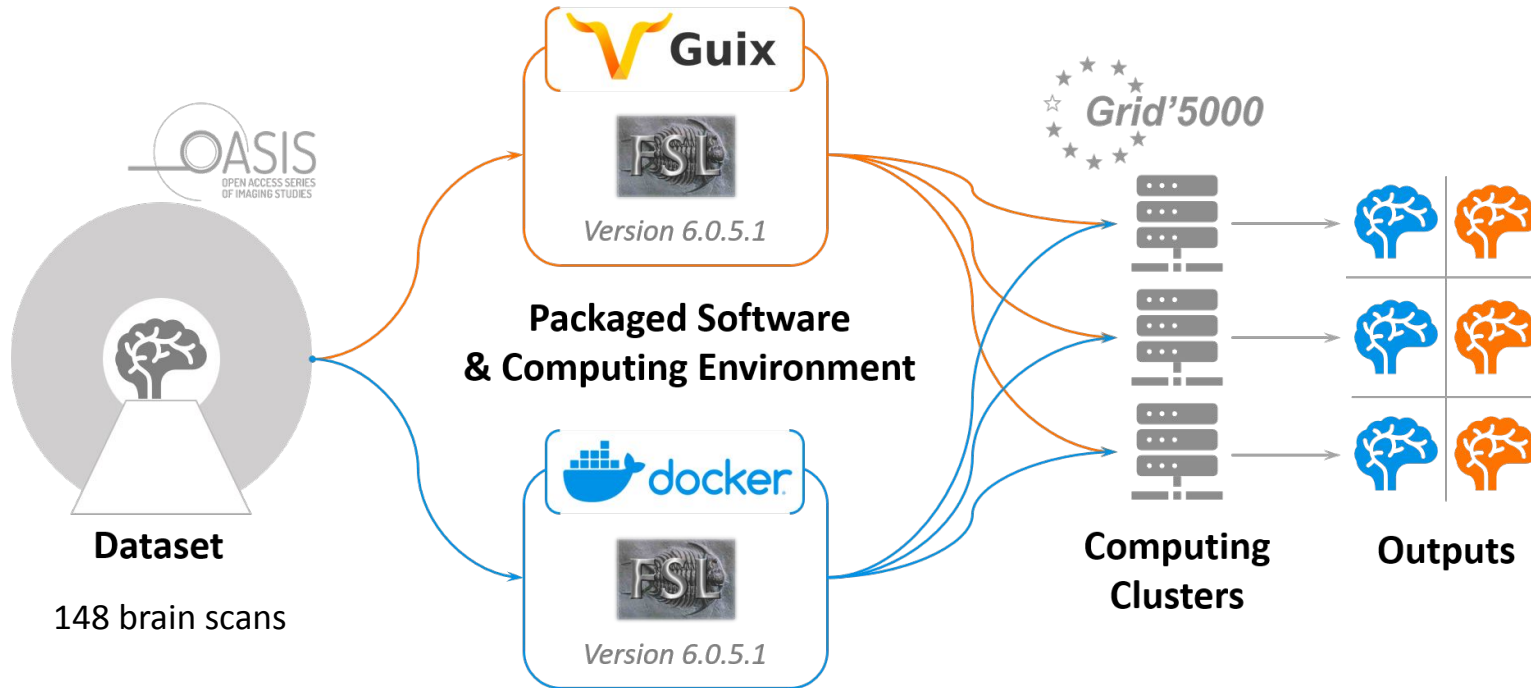

 Grid'5000 research infrastructure

- Large-scale testbed for experiment-driven research in computer science
- Access to a wide spectrum of hardware

| Cluster | CPU | Model | Micro-arch | ISE |
|----------|-------|-----------------|--------------|---------|
| uvb | Intel | Xeon X5670 | Westmere | SSE4.2 |
| hercule | Intel | Xeon E5-2620 | Sandy Bridge | AVX |
| taurus | Intel | Xeon E5-2630 | Sandy Bridge | AVX |
| parasilo | Intel | Xeon E5-2630 v3 | Haswell | AVX2 |
| nova | Intel | Xeon E5-2620 v4 | Broadwell | AVX2 |
| chiffnot | Intel | Xeon Gold 6126 | Skylake | AVX-512 |
| chiclet | AMD | EPYC 7301 | Zen | AVX2 |
| neowise | AMD | EPYC 7642 | Zen 2 | AVX2 |
| abacus21 | AMD | EPYC 7F72 | Zen 2 | AVX2 |

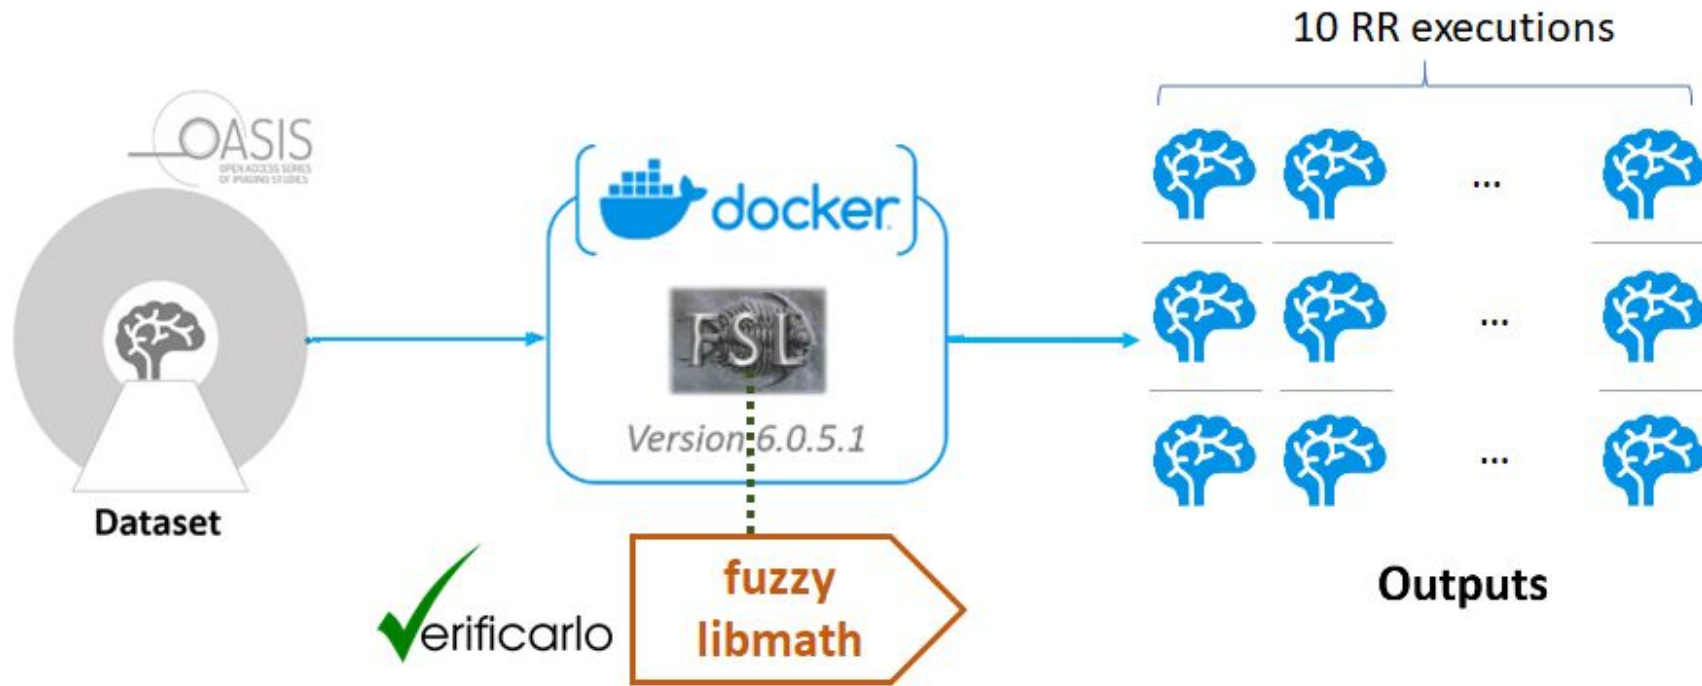
← Fused Multiply-Add (FMA)

Overview of experiments on Grid'5000



(4 Guix + 1 Docker executables) x 9 clusters = 45 experiments

Overview of MCA experiments



Hardware variability

- Comparison of global checksums
 - tarball of the 148 results for each one of the 45 experiments


| Deployment | Compilation flags (-march=) | Microarchitecture of the execution node | ISE | Global checksum |
|------------|-----------------------------|---|--------------------|-----------------|
| Docker | x86_64 | Intel Westmere, Sandy Bridge | SSE4.2, AVX | 03f... |
| Docker | x86_64 | Intel Haswell, Broadwell, Skylake, AMD Zen, Zen 2 | AVX-2 | 6a9... |
| Guix | x86_64 | All | SSE4.2, AVX, AVX-2 | b48... |
| Guix | sandybridge | Intel Sandy Bridge | AVX | b48... |
| Guix | haswell or skylake | Intel Haswell, Broadwell, Skylake, AMD Zen, Zen 2 | AVX-2 | 75e... |
| Guix | sandybridge | Intel Westmere | SSE4.2 | incompatibility |
| Guix | haswell or skylake | Intel Westmere, Sandy Bridge | SSE4.2, AVX | incompatibility |

Four different global checksums

Two micro-architecture subsets: with and without AVX-2


Variability depends on input data

- Comparison of the 148 individual results among the four sets of results
- Three of the four sets share a few identical results



Intersections between result sets (individual matrix files) for three of the four experiments.



Variability across subjects



Framewise displacement across subjects

=> importance of using large image databases

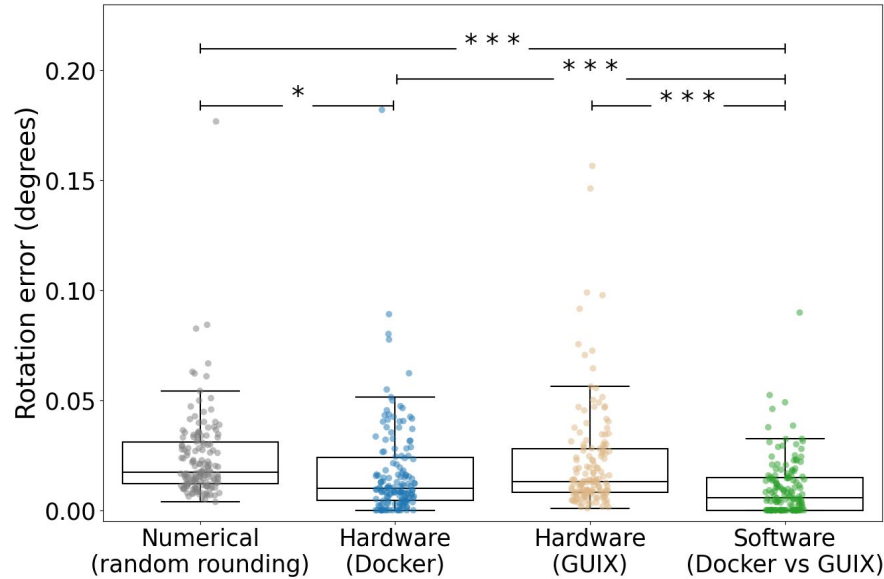
Effects on the registration



Distributions of rotation and translation differences in the transformation matrix results (‘.mat’ result files)

Differences between outputs (belonging to groups Docker-6a9 and Guix-b48) with the largest difference in translation and rotation (subject 31, scan 2)

Variabilities of comparable magnitude



Comparison between rotation errors for numerical, hardware, and software variability, for each subject

Correlation

| Variability 1 | Variability 2 | Spearman correlation |
|---------------------------|-------------------|----------------------|
| Hardware (Docker) | Numerical (RR) | 0.04 |
| Hardware (Guix) | Numerical (RR) | 0.11 |
| Software (Docker vs Guix) | Numerical (RR) | -0.11 |
| Software (Docker vs Guix) | Hardware (Guix) | 0.20 |
| Software (Docker vs Guix) | Hardware (Docker) | 0.11 |
| Hardware (Guix) | Hardware (Docker) | 0.41* |

Correlations in translation vectors

| Variability 1 | Variability 2 | Spearman correlation |
|---------------------------|-------------------|----------------------|
| Hardware (Docker) | Numerical (RR) | 0.00 |
| Hardware (Guix) | Numerical (RR) | 0.01 |
| Software (Docker vs Guix) | Numerical (RR) | -0.12 |
| Software (Docker vs Guix) | Hardware (Guix) | 0.22 |
| Software (Docker vs Guix) | Hardware (Docker) | 0.08 |
| Hardware (Guix) | Hardware (Docker) | 0.36* |

Correlations in rotation vectors

Conclusions


- Hardware, software and numerical variability lead to variations
 - of similar magnitudes but
 - uncorrelated with each other
- RR introduces perturbations of similar magnitude
 - Practical method to simulate both hardware and OS updates
- Variations remained moderate but might impact downstream analyses
- Both packaging solutions (Docker and Guix) produced
 - Each one bit-wise reproducible results when using the same packaged FLIRT executable on equivalent micro-architectures
 - Different outputs from one another due to the software variability

Discussion

- Packaging solutions
 - Docker image: little or no information on how the executable was built
 - Guix: full transparency on both compiling and runtime environments
- Compilation options
 - Only studied the impact of the “march” flag, directly related to hardware
 - Other compilation options (e.g. optimization levels) are also known to impact reproducibility
- In our case hardware variability was due to AVX-2 support
 - Further work is needed for a finer analysis of the differences observed

Availability of code and data

- FSL Docker image: hub.docker.com/r/vnmd/fsl_6.0.5.1
- FSL Guix modules and Grid'5000 scripts: gitlab.in2p3.fr/reprovip/reprovip-guix, also archived on [Software Heritage](#)
- Post-processing for the reproducibility analysis: gitlab.in2p3.fr/reprovipgroup/reprovip-notebooks-andscripts, also archived on [Software Heritage](#)
- Dockerfile, Boutiques descriptor and analysis scripts used for random rounding: github.com/big-data-lab-team/fuzzy-linreg, also archived on [Software Heritage](#)
- The OASIS dataset used in our experiments: oasisbrains.org/#data
- The transformation matrix results (.mat files): zenodo.org/records/10649569



Thank you for your attention!
Questions?

ReproVIP **anr**[®]